

1 — **Abstract** —

2 This paper presents a specification framework for monadic, recursive, interactive programs that
3 supports auto-active verification, an approach that combines user-provided guidance with automatic
4 verification techniques. This verification tool is designed to have the flexibility of a manual approach
5 to verification along with the usability benefits of automatic approaches. We accomplish this by
6 augmenting Interaction Trees, a Coq datastructure for representing effectful computations, with
7 logical quantifier events. We show that this yields a language of specifications that are easy to
8 understand, automatable, and are powerful enough to handle properties that involve non-termination.
9 Our framework is implemented as a library in Coq. We demonstrate the effectiveness of this framework
10 by verifying real, low-level code.

11 **2012 ACM Subject Classification** Theory of computation → Denotational semantics; Theory of
12 computation → Programming logic; Theory of computation → Separation logic

13 **Keywords and phrases** coinduction, specification, verification, monads

14 **Digital Object Identifier** 10.4230/LIPIcs...



© Author: Please provide a copyright holder;
licensed under Creative Commons License CC-BY 4.0



Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Interaction Tree Specifications: A Framework for Specifying Recursive, Effectful Computations that Supports Auto-active Verification

Lucas Silver ✉

University of Pennsylvania, USA

Eddy Westbrook ✉

Galois, Inc.

Matthew Yacavone ✉

Galois, Inc.

Ryan Scott ✉

Galois, Inc.

1 Introduction

Formal verification is starting to see adoption in industry as a tool for ensuring the security and correctness of software. For instance, the formally verified seL4 microkernel [13] has established a foundation that is seeing investment from a wide variety of industrial partners. Block-chain companies are using formal verification to ensure the security of cryptocurrency [15]. Amazon has even incorporated formal verification into the CI/CD process of their s2n cryptographic library [7].

Unfortunately, formal verification still remains expensive, not just in terms of time and effort but also in terms of the expertise required to formally verify a system. A number of powerful frameworks have been developed for manual formal verification, including Iris [12], VST [2], and FCSL [24]. These frameworks can specify a wide array of behaviors on a wide array of languages, but they require an expert to be used effectively. Other powerful frameworks have been developed for automatic verification, including approaches such as bounded model-checking [4] and property-directed reachability [5]. While these approaches can be operated by non-experts, they are limited in their expressiveness, leaving important properties unverified.

It is particularly difficult to reason about low-level code that contains complicated manipulations of pointer structures on the heap, as is common in languages like C, C++, and LLVM. Recently, researchers have tackled this problem using the observation that programs that are well-typed in a memory-safe, Rust-like type system are basically functional programs [3, 9, 10, 17, 18]. That is, there exists a program in a functional language whose behavior is equivalent to the original, heap-manipulating program. We call this functional program a *functional specification*. While many projects rely only implicitly on the functional specification, some, like the Heapster project [9], reify functional specifications into concrete code. Engineers can then verify properties about the derived functional code, and ensure those properties hold on the original program.

The Heapster tool consists of two components: a memory-safe type system for LLVM code, and a translation tool that produces an equivalent functional program from any well-typed LLVM program. Heapster uses these components to break verification of heap manipulating programs into two phases: a memory-safe type-checking phase that generates a monadic, recursive, interactive program that is equivalent to the original program; and a behavior-verification phase that ensures that the generated program has the correct behavior. Previous work has left open major questions about the behavior verification phase, namely, what should the language of specifications be and how do we actually prove that the programs



© Author: Please provide a copyright holder;

licensed under Creative Commons License CC-BY 4.0



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

60 satisfy the specifications.

61 This work answers these questions by developing a logic well-suited to reasoning about
62 the programs output by Heapster, as well as tools to work with these logical formulae. Taken
63 together, the Heapster tool and this work form a two-step pipeline for verifying low-level,
64 heap manipulating programs. Heapster transforms low-level, heap manipulating programs
65 into equivalent functional programs. The techniques in this paper enable proof engineers to
66 write and prove specifications over the resulting functional programs.

67 In this work, we present *interaction tree specifications*, or ITree specifications. ITree
68 specifications are an *auto-active verification framework* for *monadic, recursive, interactive*
69 *programs* based on *interaction trees* [29], or ITrees. Auto-active verification is a verification
70 technique that merges user input and automated reasoning to leverage the benefits of each.
71 Monadic, recursive, interactive programs have the ability to diverge, can interact with
72 their environment, but otherwise act as pure functional programs. Interactions with the
73 environment can include making a system call, sending a message from a server, and throwing
74 an error. ITrees are a model for monadic, recursive, interactive programs formalized in Coq.
75 ITree specifications are designed to be able to write and verify specifications about the output
76 programs of the Heapster translation tool, which are written in terms of ITrees.

77 The main body of work that takes on the task of verifying monadic programs is the
78 Dijkstra monad literature [1, 16, 27, 28]. However, most of the Dijkstra monad literature
79 cannot handle the kinds of termination sensitive specifications that we need. These papers
80 either assume a strongly normalizing language, or handle only partial specifications. The
81 exception to this is the work of Silver and Zdancewic [25]. However, while that work does
82 have a rich enough specification language for our goals, it has two significant shortcomings.
83 First, the work provides no reasoning principles for arbitrary recursive specifications. Second,
84 the work does not attempt to automate the verification of these specifications. Our work
85 accomplishes both of these goals.

86 This work is based on the idea of augmenting ITrees with operations for logical quantifiers.
87 We show that this idea leads to a language of specifications that is:

- 88 ■ easy to read, because the specifications are simply programs annotated with logical
- 89 quantifiers,
- 90 ■ capable of encoding recursive specifications, because the underlying computational lan-
- 91 guage has a powerful recursion operator, and
- 92 ■ amenable to auto-active verification, because specifications are syntactic constructs
- 93 enabling syntax-directed inference rules.

94 ITrees represent computations as potentially infinite trees whose nodes are labelled with
95 *events*. Events are syntactic representations of computational effects, like raising an error,
96 or sending data from a server. ITrees can be used to represent the semantics of recursive,
97 monadic, interactive programs. ITree specifications are ITrees enriched with events for logical
98 quantifiers. This language of specifications has the capability to express purely executable
99 computations, fully abstract specifications, and combinations of both. For example, consider
100 the following executable specification `server_impl` for a simple server program that sorts lists
101 which are sent to it:

```
102 Definition server_impl : unit → itree_spec E void :=
103   rec_fix_spec (fun rec _ =>
104     l ← trigger rcvE;;
105     ls ← sort l;;
106     trigger (sendE ls);;
107     rec tt
108   ).
109
```

XX:4 Interaction Tree Specifications

```
Class EncodingType (E:Type) : Type :=  
  response_type : E → Type.
```

■ **Figure 1** EncodingType typeclass definition

111 This specification is defined with `rec_fix_spec`, a recursion operator (defined in Section 4)
112 where applications of the `rec` argument correspond to recursive calls. The body of the recursive
113 function first calls `trigger rcvE`, which triggers the use of the receive event `rcvE`, causing the
114 program to wait to receive data. The list `l` that is received is then passed to the `sort` function,
115 defined in Section 6, which is a recursive implementation of the merge sort algorithm. Finally,
116 the sorted list returned by `sort` is sent as a response with `trigger (sendE ls)`, and the server
117 program loops back to the beginning by calling `rec`.

118 Now, consider the following specification of the behavior of our server using a combination
119 of executable and abstract features:

```
120 Definition server_spec : unit → itree_spec E void :=  
121   rec_fix_spec (fun rec _ =>  
122     l ← trigger rcvE;;  
123     ls ← exists_spec (list nat);;  
124     assert_spec (Permutation l ls);;  
125     assert_spec (sorted ls);;  
126     trigger (sendE ls);;  
127     rec tt).  
128
```

130 This function acts mostly like `server_impl` but, instead of computing a sorted list, it uses the
131 existential quantification operation `exists_spec` to introduce the list value `ls`, which it then
132 asserts is a sorted permutation of the initial list. By leaving this part of the specification
133 abstract, it allows the user to express that it is unimportant how the list is sorted, as long as
134 the response is a sorted permutation of the input list. The send and receive events, however,
135 are left concrete, allowing the user to specify what monadic events should be triggered in
136 what order. This specification implicitly defines a liveness property of the server, it will
137 reject any program that fails to eventually perform the next send or receive. By using a
138 single language for programs and specifications, our approach provides a natural way for
139 users to control how concrete or abstract the various portions of their specifications are. Our
140 approach then provides auto-active tools for proving that programs refine these specifications.

141 Necessary background explaining ITrees and Heapster is given in Section 2 and Section 3.
142 The contributions of this paper are as follows:

- 143 ■ ITree specifications, a data structure for representing specifications over monadic, recursive,
144 interactive programs, presented in Section 4
- 145 ■ a specification refinement relation over ITree specifications, along with collection of
146 verified, syntax-directed proof rules for refinement also presented in Section 4,
- 147 ■ tools for encoding and proving refinements involving total correctness specifications in
148 ITree specifications presented in Section 5,
- 149 ■ an auto-active verification technique briefly discussed in Section 6
- 150 ■ an evaluation of the presented techniques in the form of verifying a collection of realistic
151 C functions using ITree specifications and Heapster presented in Section 6.

152 2 Background

153 ITrees are a formalization for denotational semantics implemented as a coinductive variant of
154 the free monad in Coq. ITrees represent programs as potentially infinite trees. The nodes of

155 these trees are labelled with *events*. Events can, depending on the context, either represent
 156 algebraic effects or recursive function calls. The `ITree` type is parameterized by a return
 157 type `R` and a type family `E`, where `E` has an instance of the `EncodingType` type class defined in
 158 Figure 1. The `EncodingType` type class consists of function, named `response_type`, from `E` to
 159 `Type`. A value of type `itree E R` is a potentially infinite tree whose internal nodes are each
 160 labelled with an *event* `e` of type `E`, with one branch for each element of the `response_type e`
 161 whose leaves are labelled with an element of type `R`. Such a tree represents an effectful
 162 computation, where the leaves represent termination of the computation with a return value
 163 in `R` while the nodes represent uses of monadic effects. The event `e` of type `E` that labels
 164 a node represents a monadic effect that returns a value of type `response_type e`, and the
 165 children of that node represent the possible continuations of that computation depending on
 166 the return value of the effect. This is formalized in the following Coq code¹.

```
167   CoInductive itree (E : Type) {EncodingType E} (R : Type) :=
168     | Ret (r : R)
169     | Tau (t : itree E R)
170     | Vis (e : E) (k : response_type e → itree E R).
```

173 The `ITree` datatype has three constructors. The `Ret` constructor represents a pure computation
 174 that simply returns a value. The `Ret` constructor forms the leaves of an `ITree`. The `Tau`
 175 constructor represents one step of silent internal computation followed by another `ITree`.
 176 Finally, the `Vis` constructor contains an event `e` along with a continuation function `k` which
 177 defines all the branches of this `Vis` node.

178 Because `ITrees` are defined coinductively, we can construct `ITrees` with infinitely long
 179 branches. Such `ITrees` represent divergent computations. For example, the following code
 180 describes an `ITree` that consists of an infinite stream of `Tau` constructors with no events.

```
181   CoFixpoint spin : itree E R := Tau spin.
```

184 In practice, `ITrees` often end up using an event type family `E` that is a composition of
 185 several smaller type families combined in a large sum. This can easily clutter and complicate
 186 the notation. To avoid this burden, the `ITrees` library introduces the `ReSum` typeclass defined
 187 in Figure 2. An instance of `ReSum E1 E2`, written `E1 -< E2`, contains two functions: the
 188 `resum` function that injects an element of `E1` into `E2`, and the `resum_ret` function that maps
 189 elements from the response type of `resum e` to the response type of `e`. It can be thought
 190 of as a kind of subevent typeclass. The `ReSum` typeclass allows for the definition of the
 191 `trigger` function in Figure 2. The `trigger` function takes an event `e : E1` and injects it into
 192 `itree E (response_type e)` by injecting `e` into `E2`, placing that in a `Vis` node, and applying
 193 the `resum_ret` function to the response.

194

195 2.1 Equivalence up to Tau

196 One of the major advantages of the `ITrees` library is its rich equational theory. The primary
 197 notion of equivalence used for `ITrees` is called `eut` or *equivalence up to tau*. Xia et al. [29]
 198 defines `eut` as a bisimulation relation that quotients out finite differences in the number
 199 of `Tau` constructors. We use this relation because `Tau` constructors are supposed to indicate
 200 *silent* steps of computation. Ignoring finite numbers of `Tau` constructors lets us equate two
 201 `ITrees` that vary only in the number of silent computation steps.

¹ In the actual formalization, we use a negative coinductive types presentation of this data structure.

XX:6 Interaction Tree Specifications

```
Class ReSum (E1 : Type) (E2 : Type) `{EncodingType E1} `{EncodingType E2} :=
{
  resum : E1 → E2;
  resum_ret : forall {e : E1}, response_type (resum e) → response_type e;
}.

```

Notation "E1 -< E2" := (ReSum E1 E2) (at level 10).

```
Definition trigger {E1 E2} `{EncodingType E1} `{EncodingType E2} `{E1 -< E2} :forall (e1
  : E1), (itree E2 (response_type e1)) :=
  fun e ⇒ Vis (resum e) (fun x ⇒ Ret (resum_ret x)).

```

■ Figure 2 ReSum Definition

Example spin \approx spin.

Example Tau (Ret 0) \approx Ret 0.

Example \sim (spin \approx Ret 0).

■ Figure 3 eutt Examples

202 The `eutt` relation is parameterized by a relation `RR` over return values. If the relation
203 `RR` is *heterogeneous*, relating values over distinct types `R1` and `R2`, then `eutt RR` is also a
204 heterogeneous relation over `itree E R1` and `itree E R2`. Intuitively, if `eutt RR t1 t2`, then the
205 `Vis` nodes of `t1` precisely match those of `t2`, and if equivalent paths in `t1` and `t2` lead to the
206 leaves `Ret r1` and `Ret r2` then the values `r1` and `r2` are related by `RR`. Often, we are interested
207 in `eutt eq` and denote this relation with the symbol \approx .

208 The `eutt` relation is implemented in Coq using both *inductive* and *coinductive* techniques.
209 Observe the following definition of `eutt`:

```
210 Inductive euttF (RR : R1 → R2 → Prop) (sim : itree E R1 → itree E R2 → Prop) :
211   itree E R1 → itree E R2 → Prop :=
212   | eutt_Ret (r1 : R1) (r2 : R2) : euttF RR sim (Ret r1) (Ret r2)
213   | eutt_Tau (t1 : itree E R1) (t2 : itree E R2) :
214     sim t1 t2 → euttF RR sim (Tau t1) (Tau t2)
215   | eutt_Vis (e : E) (k1 : response_type e → itree E R1)
216     (k2 : response_type e → itree E R2) :
217     (forall a, sim (k1 a) (k2 a)) → euttF RR sim (Vis e k1) (Vis e k2)
218   | eutt_TauL (t1 : itree E R1) (t2 : itree E R2) :
219     euttF RR sim t1 t2 → euttF RR sim (Tau t1) t2
220   | eutt_TauR (t1 : itree E R1) (t2 : itree E R2) :
221     euttF RR sim t1 t2 → euttF RR sim t1 (Tau t2).
222

```

```
224 Definition eutt (RR : R1 → R2 → Prop) := gfp (euttF RR).
225

```

227 The `euttF` relation is an inductively defined relation, defined in terms of the `sim` argument.
228 The `eutt` relation is then defined as the greatest fixpoint of `euttF`. In this paper, all greatest
229 fixpoints are defined using the `paco` library[11] for coinductive proofs. Calls to the `sim`
230 argument in the definition of `euttF` correspond to coinductive calls of `eutt`. Recursive calls
231 to `euttF` correspond to inductive calls of `eutt`. This method of defining `eutt` allows the
232 coinductive constructors to be called infinitely often in sequence, while only a finite number
233 of calls to inductive constructors can be called without an intervening call to a coinductive
234 constructor. Specifically, only finitely many `eutt_TauL` and `eutt_TauR` steps, that remove a
235 `Tau` from only one side, are allowed before one of the remaining rules must be used to relate
236 the same constructor on both sides.

237 This definition allows us to achieve our goal of ignoring any finite difference in numbers
238 of `Tau` constructors. In particular the equations and inequalities presented in Figure 3 hold.

239 ITrees form a monad. Monads are type families with a `ret` combinator that denotes a
240 pure value, and a `bind` combinator that sequentially composes two monadic computations
241 into one. The `ret` combinator is implemented with the `Ret` constructor, while the `bind t k`
242 combinator is implemented as a coinductive function that traverses the ITree `t` and replaces
243 each leaf `Ret r` with the new subtree `k r`. This is implemented in the following Coq code:

```
244   CoFixpoint bind (t : itree E R) (k : R → itree E S) :=
245     match t with
246     | Ret r ⇒ k r
247     | Tau t ⇒ Tau (bind t k)
248     | Vis e kvis ⇒ Vis e (fun x ⇒ bind (kvis x) k)
249     end.
250
```

252 2.2 Mutually Recursive Computations

253 This section explains the recursion operator introduced by Xia et al. [29]. That work
254 demonstrated how to use events as a piece of syntax for writing collections of mutually
255 recursive functions over ITrees. Specifically, it introduced the `mrec` combinator, which lifts
256 a collection of function bodies that syntactically reference one another to a collection of
257 actually recursive functions. A similar recursion combinator is used extensively in Section 4
258 and Section 6.

259 When using the `mrec` combinator, you must first choose an event type `D`, with an
260 `EncodingType` instance, to serve as the type of recursive calls. An element `d : D` packages
261 together the choice of the function being called along with the arguments being supplied
262 to that function. The return type of the function call `d` is `response_type d`. In this context,
263 an ITree with the type `itree (D + E) R` represents the body of a mutually recursive function
264 viewing the recursive calls as inert `D` events. This ITree defines a recursive function in terms of
265 *syntactic* recursive calls. In order to resolve these syntactic recursive calls, we need a mapping
266 from recursive calls to a single layer of unfolding of the recursive function. This is represented
267 as a function of type `bodies : forall (d:D), itree (D + E) (response_type d)`. The variable
268 name `bodies` refers to the fact that this term represents the body of each function in this
269 collection of mutually recursive functions. We can then take this ITree, corecursively replace
270 each `d : D` event with the unfolded function body `bodies d`, and then repeat the process with
271 the resulting ITree. This is formalized in the following `interp_mrec` function.

```
272   CoFixpoint interp_mrec {R : Type}
273     (bodies : forall (d:D), itree (D + E) (response_type d))
274     (t : itree (D + E) R) : itree E R :=
275     match t with
276     | Ret r ⇒ Ret r
277     | Tau t ⇒ Tau (interp_mrec bodies t)
278     | Vis (inr e) k ⇒ Vis e (fun x ⇒ interp_mrec bodies (k x))
279     | Vis (inl d) k ⇒ Tau (interp_mrec bodies (bind (bodies d) k))
280     end.
281
```

283 Given this function that can resolve the recursive calls in an ITree, we can define the `mrec`
284 function that takes an initial recursive call `init : D` and computes its result.

```
285   Definition mrec (bodies : forall (d:D), itree (D + E) (response_type d)) (init : D)
286     :=
287     interp_mrec bodies (bodies init).
288
```

290 Figure 4 provides an example of a mutually recursive function defined with `mrec`. The
291 `evenoddE` type represents calls to compute the parity of a natural number. The `evenodd`
292 function computes either the `even` or the `odd` function depending on the initial recursive call

XX:8 Interaction Tree Specifications

```
Variant evenoddE : Type:=
| even (n : nat) : evenoddE
| odd (n : nat) : evenoddE.
Instance EncodingType_evenoddE : EncodingType evenoddE := fun _ => bool.

Definition evenodd_body : forall eo : evenoddE, (itree (evenoddE + voidE)) (
  response_type eo) :=
fun eo =>
  match eo with
  | even n => if Nat.eqb n 0
    then Ret true
    else trigger (odd (n -1))
  | odd n => if Nat.eqb n 0
    then Ret false
    else trigger (even (n -1))
  end.
Definition evenodd : evenoddE → itree voidE bool :=
mrec evenodd_body.
```

■ Figure 4 evenodd Definition

```
Definition Rel (A B : Type) : Type := A → B → Prop.
Definition PostRel (D1 D2 : Type) {EncodingType D1} {EncodingType D2} : Type :=
  forall (d1 : D1) (d2 : D2), response_type d1 → response_type d2 → Prop.

Inductive RComposePostRel
(R1 : Rel D1 D2) (R2 : Rel D2 D3) (PR1 : PostRel D1 D2) (PR2 : PostRel D2 D3) :
PostRel D1 D3 :=
| RComposePostRel_intros (d1 : D1) (d3 : D3) (a : response_type d1) (c :
  response_type d3) :
(forall (d2 : D2), R1 d1 d2 → R2 d2 d3 →
  exists b, PR1 d1 d2 a b ∧ PR2 d2 d3 b c) →
RComposePostRel R1 R2 PR1 PR2 d1 d3 a c.
```

■ Figure 5 Heterogeneous Event Relation Types

293 event that it is given. The `evenodd` function defines these computations mutually recursively
294 using the `mrec` function.

295 This section briefly introduces the classes of relations that we will need in order to reason
296 about specification refinement in the presence of mutually recursive computations. The
297 definition of `eut` is parameterized by a return relation, making it easy to define a relation for
298 ITrees that have identical tree structures up to `Taus`, with identical event nodes, but allows
299 freedom to choose what conditions to enforce on return values. It is natural to consider
300 generalizing `eut` to allow variation not only in the return values but also in the event nodes.
301 This kind of generalization is explored in Silver and Zdancewic [25]². The generalized relation
302 analyzes uninterpreted events, typically those representing recursive function calls, with
303 respect to pre-conditions and post-conditions. We want to relate `Vis` nodes whose events
304 satisfy the pre-condition and whose continuations are related given any inputs that satisfy
305 the post-condition. This corresponds to assuming that two function calls return related
306 outputs as long as they are given related inputs.

307 Definitions of pre-condition and post-condition types are presented in Figure 5. Pre-
308 conditions, `Rel`, are encoded as two-argument, heterogeneous relations, i.e. functions of type
309 `D → E → Prop`, and utilize standard relational combinators like relational sums, `sum_rel`, and

² In Silver and Zdancewic [25] this relation is referred to as `eutEv`. It has since been renamed to `rut` in release branches of the Interaction Trees library.

310 relational composition, `rcompose`. Post-conditions, `PostRel`, are encoded as four-argument, de-
 311 pendent relations. In particular, `forall (d:D) (e:E), encoded_by d → encoded_by e → Prop`,
 312 where both `D` and `E` have an `EncodingType` instance. Intuitively, post-conditions are a function
 313 from events to relations over their response types. These post-conditions admit a standard
 314 definition of relational sums. For relational composition, in addition to requiring two `PostRel`
 315 relations, it also requires two standard relations, called *coordinating relations*. The full
 316 definition is presented in Figure 5.

317 To relate four values `d1:D1`, `d3:D3`, `a:encoded_by d1`, `c:encoded_by d3`, we require that
 318 given any `d2:D2` that is related by the coordinating relations to `d1` and `d3`, there exists a
 319 `b:encoded_by d2` such that both `PR3 d1 d2 a b` and `PR4 d2 d3 b c`.

320 Later in the paper, we recover an eutt-like definition of specification refinement by
 321 specializing the event relations to be an appropriate form of equality. For `Rel`, this is precisely
 322 the equality relation. For `PostRel`, we define an inductive datatype that enforces equality on
 323 response values.

```
324 Variant PostRelEq : PostRel E E :=
325   PostRelEq_intro e a : PostRelEq e e a a.
```

328 **3 Specification Extraction with Heapster**

329 This section introduces the Heapster tool for specification extraction. We present Heapster
 330 in order to provide context for the evaluation of this work in Section 6. In the evaluation, we
 331 demonstrate how effective ITree specifications can be when paired with a tool like Heapster.
 332 We start with a collection of low-level, heap manipulating C programs, use Heapster to
 333 produce equivalent functional programs, and finally use ITree specifications to specify and
 334 verify the output programs.

335 There is a growing body of work [3, 9, 17, 18] based on the idea that programs that
 336 satisfy memory-safe type systems like Rust can be represented with equivalent functional
 337 programs. Rust’s pointer discipline, which ensures that all pointers in a program are either
 338 shared read or exclusive write, allows us to reason about the effects of pointer updates purely
 339 locally. This locality property can be used to define a pure functional model, referred to as a
 340 *functional specification*, of the behaviors of a program, which can in turn be used to verify
 341 properties of that program.

342 Whereas some work uses this notion of a functional model implicitly, *specification ex-*
 343 *traction* is the idea that the functional model can be extracted automatically as an artifact
 344 that can be used for verification. Specification extraction separates verification into two
 345 phases: a type-checking phase, where the functions in a program are type-checked against
 346 user-specified memory-safe types; and a behavior verification phase, where the user verifies
 347 the specifications that are extracted from this type-checking process. The Heapster tool[9] is
 348 an implementation of the idea of specification extraction. Heapster provides a memory-safe,
 349 Rust-like type system for LLVM, along with a typechecker. Heapster also provides a transla-
 350 tion from well-typed LLVM programs to monadic, recursive, interactive programs, modeled
 351 with ITrees, that describe a behavioral model of the original program. This translation is
 352 inspired by the Curry-Howard isomorphism. Heapster types are essentially a form of logical
 353 propositions regarding the heap, so, by the Curry-Howard isomorphism, it is natural to view
 354 typing derivations, a form of proof, as a program. We give a brief overview of the Heapster
 355 type system and its specification extraction process in this section and illustrate it with an
 356 example.

XX:10 Interaction Tree Specifications

Value Types	T	$::=$	$\text{bv } n \mid \text{llvmptr } n \mid \dots$
Expressions	e	$::=$	$n \mid \text{llvmword } e \mid \dots$
RW Modality	rw	$::=$	$W \mid R$
Permissions	τ	$::=$	$\text{ptr}((rw, e) \mapsto \tau) \mid \tau_1 * \tau_2 \mid \tau_1 \vee \tau_2 \mid \exists x:T.\tau \mid \text{eq}(e) \mid \mu X.\tau \mid X \mid \dots$

■ **Figure 6** An Abbreviated Grammar of the Heapster Type System

357 The Heapster type system is a permission type system. Typing assertions of the form
358 $x : \tau$ mean that the current function holds permissions to perform actions allowed by τ
359 on the value contained in variable x . The central permission construct of Heapster is the
360 permission to read or write a pointer value. Like Rust, Heapster is an affine type system,
361 meaning that the permissions held by a function can change at different points in the function.
362 In particular, a command can consume a permission, preventing further commands from
363 using that permission again. Also like Rust, Heapster allows read-only permissions to be
364 duplicated, allowing multiple read-only pointers to the same address, but does not allow
365 write permissions to be duplicated. This enforces the invariant that all pointers are either
366 shared read or exclusive write, a powerful property for proving memory-safety.

367 Figure 6 gives an abbreviated grammar for the Heapster type system. The value types T
368 are inhabited by pieces of first order data. In particular, they contain the type $\text{bv } n$ of n -bit
369 bitvectors (i.e., n -bit binary values) and the type $\text{llvmptr } n$ of n -bit LLVM values, among
370 other value types not discussed here. Heapster uses the CompCert memory model [14],
371 where LLVM values are either a word value or a pointer value represented as a pair of a
372 memory region plus an offset in that region. The expressions e include numeric literals n and
373 applications of the llvmword constructor of the LLVM value type to build an LLVM value
374 from a word value.

375 The first permission type in Figure 6, $\text{ptr}((rw, e) \mapsto \tau)$, represents a permission to read
376 or write (depending on rw) a pointer at offset e . Write permission always includes read
377 permission. This permission also gives permission τ to whatever value is currently pointed
378 to by the pointer with this permission. Permission type $\tau_1 * \tau_2$ is the separating conjunction
379 of τ_1 and τ_2 , giving all of the permissions granted by τ_1 or τ_2 , where τ_1 and τ_2 contain no
380 overlapping permissions. Permission type $\tau_1 \vee \tau_2$ is the disjunction of τ_1 and τ_2 , which either
381 grants permissions τ_1 or τ_2 . The existential permission $\exists x:T.\tau$ gives permission τ for some
382 value x of value type T . The equality permission $\text{eq}(e)$ states that a value is known to be
383 equal to an expression e . This can be viewed as a permission to assume the given value
384 equals e . Finally, $\mu X.\tau$ is the least fixed-point permission, where permission variable X is
385 bound in τ . This satisfies the fixed-point property, that $\mu X.\tau$ is equivalent to $[\mu X.\tau/X]\tau$.

386 As a simple example, the user can define the Heapster type

```
387 int64 =  $\exists x:\text{bv } 64.\text{eq}(\text{llvmword } x)$ 
```

388 This Heapster type describes an LLVM word value, i.e., an LLVM value that equals $\text{llvmword } x$
389 for some bitvector x .

390 As a slightly more involved example, consider the following definition of a linked list
391 structure in C:

```
392 typedef struct list64_t { int64_t data;  
393                          struct list64_t *next; } list64_t;  
394  
395
```

```

int64_t is_elem (int64_t x, list64_t *l) {
  x:int64,l:list64⟨R⟩
  x:int64,l:eq(llvmword 0) OR x:int64,l:ptr((R,0) ↦ int64) * ptr((R,8) ↦ list64⟨R⟩)
  if (l == NULL) {
    x:int64,l:eq(llvmword 0)
    return 0;
  } else {
    x:int64,l:ptr((R,0) ↦ int64) * ptr((R,8) ↦ list64⟨R⟩)
    if (l->data == x) { return 1; }
    else {
      list64_t *l2 = l->next;
      x:int64,l:ptr((R,0) ↦ int64) * ptr((R,8) ↦ eq(l2)),l2:list64⟨R⟩
      return is_elem (x, l2);
    }
  }
}

```

■ **Figure 7** Type-checking the `is_elem` Function Against Type $x:\text{int64}, l:\text{list64}\langle R \rangle \multimap r:\text{int64}$

396 A C value of type `list64_t*` represents a list, where a NULL pointer represents the empty list
 397 and a non-NULL pointer to a `list64_t` struct represents a list whose head is the 64-integer
 398 contained in the `data` field and whose tail is given by the `next` field.

399 The following Heapster type describes this linked list structure:

400 $\text{list64}\langle rw \rangle = \mu X. \text{eq}(\text{llvmword } 0) \vee (\text{ptr}((rw, 0) \mapsto \text{int64}) * \text{ptr}((rw, 8) \mapsto X))$

401 The $\text{list64}\langle rw \rangle$ type is parameterized by a read-write modality rw , which says whether it
 402 describes a read-only or read-write pointer to a linked list. The permission states that the
 403 value it applies to either equals the NULL pointer, represented as `llvmword 0`, or points at
 404 offset 0 to a 64-bit integer and at offset 8^3 to an LLVM value that itself recursively satisfies
 405 the $\text{list64}\langle rw \rangle$ permission. Note that the fact that it is a least fixed-point implicitly requires
 406 the list to be loop-free.

407 Figure 7 illustrates the process of Heapster type-checking on a simple function `is_elem`
 408 that checks if 64-bit integer `x` is in the linked list `l`. Note that Heapster in fact operates
 409 on the LLVM code that results from compiling this C code, but the type-checking is easier
 410 to visualize on the C code rather than looking at its corresponding LLVM. Ignoring the
 411 Heapster types for the moment, which are displayed with a grey background in the figure,
 412 `is_elem` first checks if `l` is NULL, and if so returns 0 to indicate that the check has failed. If
 413 not, it checks if the head of the list in `l->data` equals `x`, and if so, returns 1. Otherwise, it
 414 recurses on the tail `l->next`.

415 The Heapster permissions for this function are

416 $x:\text{int64}, l:\text{list64}\langle R \rangle \multimap r:\text{int64}$

417 The lollipop symbol, \multimap , is used to write Heapster function types. This type means that
 418 input `x` is a 64-bit integer and `l` is a read-only linked list pointer and the return value r is a
 419 64-bit integer value.

420 To type-check `is_elem`, Heapster starts by assuming the input types for the arguments.
 421 This is displayed in the first grey box of Figure 7. In order to type-check the NULL comparison

³ We assume a 64-bit architecture, so offset 8 references the second value of a C struct.

XX:12 Interaction Tree Specifications

422 on `l`, Heapster must first unfold the recursive permission on `l` and then eliminate the resulting
 423 disjunctive permission. This latter step results in Heapster type-checking the remaining
 424 code twice, once for each branch of the disjunct. More specifically, the remaining code is
 425 type-checked once under the assumption that `l` equals `NULL` and once under the assumption
 426 that it points to a valid `list64_t` struct. In the first case, the `NULL` check is guaranteed
 427 to succeed, and so the `if` branch is taken with those permissions, while in the second, the
 428 `NULL` check is guaranteed to fail, so the `else` branch is taken.

429 In the `if` branch, the value `0` is returned. Heapster determines that this value satisfies the
 430 required output permission `int64`. In the `else` branch, `l->data` is read, by dereferencing `l`
 431 at offset `0`. This is allowed by the permissions on `l` at this point in the code. If the resulting
 432 value equals `x`, then `1` is returned, which also satisfies the output permission `int64`. Otherwise,
 433 `l->next` is read, by dereferencing `l` at offset `0`, and the result is assigned to local variable
 434 `l2`. This assigns `list64(R)` permission to `l2`. The permission on offset `8` of `l` is updated to
 435 indicate that the value currently stored there equals `l2`. The `list64(R)` permission on `l2` is
 436 then used to type-check the subsequent recursive call to `is_elem`.

437 Once a function is type-checked, Heapster performs specification extraction to extract a
 438 pure functional specification of the function's behavior. Specification extraction translates
 439 permission types to Coq types and typing derivations to Coq programs. The type translation
 440 is defined as follows:

$$\begin{array}{l}
 \llbracket \text{ptr}((rw, e) \mapsto \tau) \rrbracket = \llbracket \tau \rrbracket \\
 \llbracket \tau_1 \vee \tau_2 \rrbracket = \llbracket \tau_1 \rrbracket + \llbracket \tau_2 \rrbracket \\
 \llbracket \text{eq}(e) \rrbracket = \text{unit} \\
 \llbracket \tau_1 * \tau_2 \rrbracket = \llbracket \tau_1 \rrbracket * \llbracket \tau_2 \rrbracket \\
 \llbracket \exists x : T. \tau \rrbracket = \{x : \llbracket T \rrbracket \ \& \ \llbracket \tau \rrbracket\} \\
 \llbracket \mu X. \tau \rrbracket = \text{user-specified type } A \\
 \text{isomorphic to } \llbracket \llbracket \mu X. \tau / X \rrbracket \rrbracket
 \end{array}$$

442 Pointer permissions `ptr((rw, e) ↦ τ)` are translated to the result of translating the permission
 443 `τ` of the value that is pointed to. This means that specification extraction erases pointer
 444 types, which are no longer needed in the resulting functional code. Conjunctive permissions are
 445 translated to pairs, disjunctive permissions are translated to sums, and existential permissions
 446 are translated to dependent pairs (using a straightforward translation $\llbracket T \rrbracket$ of value types that
 447 we omit here). The equality type `eq(e)` is translated to the Coq unit type `unit`, meaning that
 448 they contain no data in the extracted specifications. We already proved the equality in the
 449 typechecking phase, and we have no use for the particular equality proof the typechecker
 450 provided. To translate a least fixed-point type `μX.τ`, the user specifies a type that satisfies
 451 the fixed-point equation, meaning a pair of functions

$$\text{fold} : \llbracket \llbracket \mu X. \tau / X \rrbracket \rrbracket \rightarrow \llbracket \mu X. \tau \rrbracket \quad \text{unfold} : \llbracket \mu X. \tau \rrbracket \rightarrow \llbracket \llbracket \mu X. \tau / X \rrbracket \rrbracket$$

453 that form an isomorphism.

454 As an example, the translation of `int64` is the Coq sigma type `{x:bitvector 64 & unit}`.
 455 Note that Heapster will in fact optimize away the unnecessary `unit` type, yielding the type
 456 `bitvector 64`. As a slightly more complex example, in order to translate the `list64(rw)`
 457 described above, the user must provide a type `T` that is isomorphic to the type

$$\text{unit} + (\text{bitvector } 64 * T)$$

461 The simplest choice for `T` is the type `list (bitvector 64)`. In this way, the imperative
 462 linked list data structure defined above in `C` is translated to the pure functional list type.

463 Rather than defining the translation of Heapster typing derivations into Coq programs
 464 here, we illustrate the high-level concepts with our example and refer the interested reader
 465 to He et al. [9] for more detail. The translation of `is_elem` is given as a Coq specification

```

Definition is_elem_spec : bitvector 64 * list (bitvector 64) →
  itree_spec E (bitvector 64) :=
  rec_fix_spec (fun rec '(x,l) ⇒
    either
      unit (bitvector 64 * list (bitvector 64)) (* input types *)
      (itree_spec _ (bitvector 64))           (* output type *)
      (fun _ ⇒ Ret (intToBv 64 0))           (* nil case *)
      (fun '(hd,tl) ⇒                          (* cons case *)
        if bvEq 64 hd x then Ret (intToBv 64 1) (* return 1 *)
        else rec (x,tl))                       (* recursive call *)
      (unfoldList l)).                          (* unfolded argument *)

```

■ **Figure 8** Extracted Specification for `is_elem`

466 `is_elem_spec` in Figure 8. At the top level, this specification uses `rec_fix_spec` to define
 467 a recursive function to match the recursive definition of `is_elem`. This binds a local variable
 468 `rec` to be used for recursive calls to the specification.

469 To understand the rest of the specification, we step through the Heapster type-checking
 470 depicted in Figure 7. The first step of that type assignment unfolds the permission type
 471 `list64⟨W⟩` on `l`. The corresponding portion of the specification is the call to `unfoldList`,
 472 which unfolds the input list `l` to a sum of a unit or the head and tail of the list. The next step
 473 of the Heapster type-checking is to eliminate the resulting disjunctive permission on `l`. The
 474 corresponding portion of the specification is a call to the `either` sum elimination function.
 475 In the left-hand case of the disjunctive elimination, the NULL test of the C program succeeds,
 476 and 0 is returned. Similarly, in the Coq specification, the `nil` case returns the 0 bitvector
 477 value.

478 In the right-hand case of the disjunctive elimination of the Heapster type-checking, the
 479 NULL test fails, and so `l` is a valid pointer to a C struct with `data` and `next` fields. This is
 480 represented by the pattern-match on the cons case in the Coq specification, yielding variables
 481 `hd` and `tl` for the head and tail of the list. The body of this case then tests whether the head
 482 equals the input variable `x`, corresponding to the `x==l->data` expression in the C program.
 483 If so, then the bitvector value 1 is returned. Otherwise, the specification performs a recursive
 484 call, passing the same value for `x` and the tail of the input list for `l`.

485 4 ITree Specifications and Refinement

486 In this paper, we introduce a specialization of the ITree data type that encodes specifications
 487 over ITrees. To do this, we take some base event type family `E`, and extend it with constructors
 488 for universal and existential quantification. This is formalized in the following definition for
 489 `SpecEvent`.

```

490 Inductive SpecEvent (E : Type) {EncodingType E} : Type :=
491   | Spec_vis (e : E) : SpecEvent E
492   | Spec_forall (A : type) : SpecEvent E
493   | Spec_exists (A : type) : SpecEvent E
494   .
495

```

497 The `Spec_vis` constructor allows you to embed a base event `e : E` into the type `SpecEvent E`.
 498 The `Spec_forall` constructor signifies universal quantification, and the `Spec_exists` constructor
 499 signifies existential quantification. For the purposes of specifying Heapster programs, we

XX:14 Interaction Tree Specifications

500 only need to quantify over a fixed grammar of first order types⁴. This includes natural
501 numbers, bit vectors, functions, products, logical propositions, and sums. We have omitted
502 the definition of the particular fixed grammar of types used in this work for space.

503 We define *ITree specifications* as the type of ITrees with a `SpecEvent` as the event type.

```
504 Definition itree_spec (E : Type)  $\{$ EncodingType E $\}$  (R : Type) :=  
505   itree (SpecEvent E) R.  
506
```

508 Because ITree specifications are actually a special kind of ITree, they inherit all the
509 useful metatheory and code defined for ITrees. In particular, we can reason about them
510 equationally with `eut`, and apply the monad functions to them.

511 4.1 ITree Specification Refinement

512 The notion that a program adheres to a specification is defined with the notion of refinement.
513 Refinement is the main judgment involved in using ITree specifications, and is for instance
514 the primary form of proof goal proved by the provided automation tool. Intuitively, the
515 logical quantifier events mean that an ITree specification represents a set of computations. A
516 fully concrete ITree specification, with no logical quantifier events, represents a singleton set,
517 while a more abstract specification might represent a larger set. The refinement relation is
518 then defined such that, if one ITree specification refines another, then the former represents a
519 subset of the latter. So, for instance, if we prove that a concrete specification refines a more
520 abstract specification, then we have shown that the singleton program in the set represented
521 by the concrete specification satisfies the specification. Note that refinement is actually a
522 coarser relation than subset; this is discussed later in Section 4.4.

523 The ITree specification refinement relation is based on the idea of refinement of logical
524 formulae with the `eut` relation. As in a sequent calculus, we can eliminate quantifiers in our
525 specification logic using quantifiers in the base logic, in this case Coq. Quantifiers on the
526 right of a refinement get eliminated to the corresponding Coq quantifiers, while quantifiers on
527 the left get eliminated to the dual of the corresponding Coq quantifier. This means that both
528 a `Spec_forall` on the right and a `Spec_exists` on the left get eliminated to a Coq `forall`. And
529 both a `Spec_exists` on the right and a `Spec_forall` on the left get eliminated to a Coq `exists`.
530 ITree specifications form a lattice with refinement serving as the preorder, `Spec_forall` acting
531 as the complete meet, and `Spec_exists` acting as the complete join. The portions of ITree
532 specifications with computational content, including the `Ret` leaves, `Spec_vis` nodes, and silent
533 `Tau` nodes, get compared as they do in the `eut` relation.

534 The ITree specification refinement relation shares many mechanical details with the
535 `eut` relation. Both are defined by taking the greatest fixed point of an inductively defined
536 relation to get a mixture of inductive and coinductive properties. Both behave identically
537 on `Tau` and `Ret` nodes. The refinement relation differs in its inductive rules for eliminating
538 logical quantifiers, and in its usage of heterogeneous event relations to enforce pre- and post-
539 conditions on `Spec_vis` events. These pre- and post- conditions are necessary in order to give
540 the refinement relation the flexibility needed to state the reasoning principle for `mrec`. The
541 initial inductively defined relation, `refinesF`, contains the following header code.

```
542 Inductive refinesF  
543   (RPre : Rel E1 E2) (RPost : PostRel E1 E2) (RR : Rel R1 R2)  
544   (sim : itree_spec E1 R1 → itree_spec E2 R2 → Prop)  
545   : itree_spec E1 R1 → itree_spec E2 R2 → Prop :=  
546
```

⁴ While we could quantify over `Type` in these definitions, this introduces universe level constraints that we prefer to avoid

548 Much like in the definition of `euttF`, the `sim` argument represents corecursive calls of the
 549 `refines` relation, and the `RR` argument is the relation used for return. Unlike in `euttF`, `refinesF`
 550 takes in arguments for a `PreRel` and a `PostRel`. These arguments are included in order to
 551 represent pre- and post- conditions on mutually recursive function bodies.

552 The `refinesF` relation has several constructors that work precisely the same as the
 553 corresponding `euttF` constructors. These constructors define the relation's behavior on `Ret`
 554 and `Tau` nodes.

```
555 | refines_Ret (r1 : R1) (r2 : R2) : RR r1 r2 → refinesF RPre RPost RR sim (Ret r1)
556   (Ret r2)
557 | refines_Tau (phi1 : itree_spec E1 R1) (phi2 : itree_spec E2 R2) : sim phi1 phi2
558   →
559   refinesF RPre RPost RR sim (Tau phi1) (Tau phi2)
560 | refines_TauL (t1 : itree_spec E1 R1) (t2 : itree_spec E2 R2) :
561   refinesF RPre RPost RR sim t1 t2 → refinesF RPre RPost RR sim (Tau t1) t2
562 | refines_TauR (t1 : itree_spec E1 R1) (t2 : itree_spec E2 R2) :
563   refinesF RPre RPost RR sim t1 t2 → refinesF RPre RPost RR sim t1 (Tau t2)
564
```

566 The constructor dealing with `Spec_vis` nodes generalizes the constructor dealing with `Vis`
 567 nodes in `euttF`. This constructor relates `Spec_vis` nodes as long as two conditions hold on
 568 the events, `e1` and `e2`, and the continuations, `k1` and `k2`. The `ITree` specifications must satisfy
 569 the precondition, by having `e1` and `e2` satisfy `RPre`. And the `ITree` specifications must satisfy
 570 the post condition by having `k1` a refine `k2` b, whenever `a` and `b` are related by `RPost e1 e2`.

```
571 | refines_Spec_vis (e1 : E1) (e2 : E2)
572   (k1 : response_type e1 → itree_spec E1 R1) (k2 : response_type e2
573     → itree_spec E2 R2) :
574   RPre e1 e2 → (forall a b, RPost e1 e2 a b → sim (k1 a) (k2 b)) →
575   refinesF RPre RPost RR sim (Vis (Spec_vis e1) k1) (Vis (Spec_vis e2) k2)
576
```

578 The added complications of this rule allow us to reason about mutually recursive functions.
 579 It ensures that related function outputs assume that function calls with arguments related
 580 by the precondition return values related by the post condition when analyzing mutually
 581 recursive functions.

582 Finally, we need constructors dealing with quantifier events. This definition uses only
 583 inductive constructors to eliminate quantifier events. We made this choice to avoid certain
 584 peculiar issues related to `ITree` specifications that consist of infinite trees of only quantifiers.
 585 Given coinductive constructors for quantifier events, we would be able to prove that such
 586 `ITree` specifications both refine and are refined by any other arbitrary `ITree` specification.
 587 That choice would cause certain `ITree` specifications to serve as both the top and bottom
 588 elements of the refinement order. This would serve as a counterexample to the transitivity of
 589 refinement, a desired property. So we chose to only use inductive constructors for quantifier
 590 events. This means that `ITree` specifications that consist of infinite trees of only quantifiers
 591 cannot be related by refinement to any other `ITree` specifications.

592 Quantifiers on the right get directly translated into Coq level quantifiers.

```
593 | refines_forallR (t : itree_spec E1 R1) (A : type) (k : response_type A →
594   itree_spec E2 R2) :
595   (forall a, refinesF RPre RPost RR sim t (k a)) →
596   refinesF RPre RPost RR sim t (Vis (Spec_forall A) k)
597 | refines_existsR (t : itree_spec E1 R1) (A : type) (k : response_type A →
598   itree_spec E2 R2) :
599   (exists a, refinesF RPre RPost RR sim t (k a)) →
600   refinesF RPre RPost RR sim t (Vis (Spec_exists A) k)
601
```

603 Quantifiers on the left get translated into their dual quantifier at the Coq level. Eliminating
 604 a `Spec_forall` on the left gives you an `exists`. Eliminating a `Spec_exists` on the left gives you
 605 an `forall`.

XX:16 Interaction Tree Specifications

```
Class CoveredType (A : Type) := {
  encoding : type; surjection : response_type encoding → A;
  surjection_correct : forall a : A, exists x, surjection x = a; }.

Definition forall_spec {E}
  `{EncodingType E}
  (A:Type) `{CoveredType A} :
  itree_spec E A :=
  Vis (Spec_forall encoding)
    (fun x ⇒ Ret (surjection x)).

Definition exists_spec {E}
  `{EncodingType E}
  (A:Type) `{CoveredType A} :
  itree_spec E A :=
  Vis (Spec_exists encoding)
    (fun x ⇒ Ret (surjection x)).

Definition assume_spec {E}
  `{EncodingType E} (P : Prop) :
  itree_spec E unit :=
  forall_spec P;; Ret tt.

Definition assert_spec {E}
  `{EncodingType E} (P : Prop) :
  itree_spec E unit :=
  exists_spec P;; Ret tt.
```

■ Figure 9 Basic Specifications

```
606 | refines_forallL (A : type) (k : response_type (Spec_forall A) → itree_spec E1 R1)
607 | (t : itree_spec E2 R2) :
608 | (exists a, refinesF RPre RPost RR sim (k a) t) →
609 | refinesF RPre RPost RR sim (Vis (Spec_forall A) k) t
610 | refines_existsL (A : type) (k : response_type (Spec_exists A) → itree_spec E1 R1)
611 | (t : itree_spec E2 R2) :
612 | (forall a, refinesF RPre RPost RR sim (k a) t) →
613 | refinesF RPre RPost RR sim (Vis (Spec_exists A) k) t
614
```

616 This `refinesF` relation is used to define the `refines` relation as follows.

```
617 Definition refines RPre RPost RR := gfp (refinesF RPre RPost RR).
618
```

620 4.2 Padded ITrees

621 Useful refinement relations should respect the `eutt` relation. When using ITrees as a denota-
622 tional semantics, `eutt` is the basis of any program equivalence relation. Equivalent programs
623 and specifications should not be observationally different according to the refinement relation.
624 However, the `refines` relation does not respect `eutt`

625 We can easily demonstrate this with the following three ITree specifications.

```
626 CoFixpoint spin : itree_spec E R := Tau spin.
627 CoFixpoint phi1 : itree_spec E R := Vis (Spec_forall t) (fun _ ⇒ Tau (phi1)).
628 CoFixpoint phi2 : itree_spec E R := Vis (Spec_forall t) (fun _ ⇒ phi2).
629
```

631 The `spin` specification represents a silently diverging computation. The `phi1` specification
632 is an infinite stream that alternates between `Spec_forall` nodes and `Tau` constructors. The
633 `phi2` specification is a similar ITree to `phi1` that just lacks the `Tau` nodes. As these ITree
634 specifications all diverge along all paths and lack any `Spec_vis` nodes, the `RPre`, `RPost`, and `RR`
635 relations that we choose do not matter. Given any choice for those relations, `spin` refines
636 `phi1` as we can use the inductive `refines_forallL` rule to get rid of the `Spec_forall` nodes,
637 allowing us to match `Tau` nodes on both trees and apply the coinductive `refines_Tau` rule.
638 This process can be extended coinductively allowing us to construct the refinement proof.
639 The `phi1` ITree specification is `eutt` to `phi2`, as the only difference between the specifications
640 is a single `Tau` node after every `Vis_forall` node. However, `spin` does not refine `phi2`, as there
641 is no coinductive constructor that we can apply in order to write a proof for these divergent


```

CoFixpoint interp_mrec_spec {R : Type}
  (bodies : forall (d:D), (itree_spec (D + E)) (response_type d)) (t : itree_spec (D + E)
    ) R) : itree_spec E R :=
  match t with
  | Ret r => Ret r
  | Tau t => Tau (interp_mrec_spec bodies t)
  | Vis (Spec_forall A) k => Vis (@Spec_forall E _ A) (fun x : response_type (Spec_forall
    A) => interp_mrec_spec bodies (k x))
  | Vis (Spec_exists A) k => Vis (@Spec_exists E _ A) (fun x => interp_mrec_spec bodies (
    k x))
  | Vis (Spec_vis (inr e)) k => Vis (Spec_vis e) (fun x => interp_mrec_spec bodies (k x))
  | Vis (Spec_vis (inl d)) k => Tau (interp_mrec_spec bodies (bind (bodies d) k))
  end.

Definition mrec_spec (bodies : forall (d:D), (itree_spec (D + E)) (response_type d)) (
  init : D) :=
  interp_mrec_spec bodies (bodies init).

```

■ **Figure 10** `mrec_spec` Definition

ITree specifications. Problems like this arise with any ITree specifications that consist of infinitely many quantifier nodes with nothing between them.

To fix this problem, we restrict our focus to a subset of ITrees that does not include ones like `phi2`. This is the set of *padded* ITrees, in which every `Vis` node must be immediately followed by a `Tau`. We formalize this with the coinductive `padded` predicate, whose definition has been omitted to save space. The refinement relation does not distinguish between different ITree specifications that are `eutt` to one another as long as they are padded. This means that can rewrite one ITree specification into another under a refinement according to `eutt` as long as both are padded.

Furthermore, it is easy to take an arbitrary ITree, and turn it into a padded ITree. That is implemented by the `pad` function, which corecursively adds a `Tau` after every `Vis` node. From here, we can focus primarily on the following definition of `padded_refines` which pads out all ITree specifications before passing them to the `refines` relation.

```

Definition padded_refines RPre RPost RR phi1 phi2 :=
  refines RPre RPost RR (pad phi1) (pad phi2).

```

In Figure 9, we introduce several simple ITree specifications that implement quantification over some types, and assumption and assertion of propositions. The `forall_spec` and `exists_spec` specifications rely on the `CoveredType` type class. A `CoveredType` instance for a type `A` contains an element of the restricted type grammar, `encoding`, whose interpretation corresponds to `A`. It also contains a valid surjection from the interpreted type `response_type encoding` to the original type `A`. In practice, we always instantiate this surjection with the identity function, but this type class formalization gives us the tools that we need without needing to do too much dependently typed programming. We can use `forall_spec` and `exists_spec` to define assumption and assertion, respectively, as `Prop` is part of the restricted grammar of types that `SpecEvent` can quantify over.

4.3 Padded Refinement Meta Theory

This subsection introduces some of the useful, verified metatheory we provide for ITree specifications in terms of `padded_refines` relation.

We prove that we can compose refinement results with the monadic `bind` operator.

XX:18 Interaction Tree Specifications

```
674 Theorem padded_refines_bind (phi1 : itree_spec E1 R1) (phi2 : itree_spec E2 R2)
675   (kphi1 : R1 → itree_spec E1 S1)
676   (kphi2 : R2 → itree_spec E2 S2) :
677   padded_refines RPre RPost RR phi1 phi2 →
678   (forall r1 r2, RR r1 r2 → padded_refines RPre RPost RS (kphi1 r1) (kphi2 r2)) →
679   padded_refines RPre RPost RS (bind phi1 kphi1) (bind phi2 kphi2).
680
681
```

682 We prove that the `padded_refines` relation is transitive. To state the transitivity result in
683 full generality, we need to use the composition relation introduced in Figure 5.

```
684 Theorem padded_refines_trans : forall (phi1 : itree_spec E1 R1) (phi2 : itree_spec E2
685   R2) (phi3 : itree_spec E3 R3),
686   padded_refines RPre1 RPost1 RR1 phi1 phi2 →
687   padded_refines RPre2 RPost2 RR2 phi2 phi3 →
688   padded_refines (RCompose RPre1 RPre2)
689   (RComposePostRel RPost1 RPost2) (RCompose RR1 RR2) phi1 phi3.
690
691
```

692 We prove a reasoning principle for mutually recursive specifications as well. To do
693 this, we first provide a slightly different definition of mutual recursion that handles the
694 quantifier events correctly, defined in Figure 10. The key to proving refinements between
695 `mrec_spec` specifications is to use the `PreRel` and `PostRel` relations to establish pre- and post-
696 conditions on recursive calls. This involves choosing a `PreRel` over recursive call events,
697 `RPreInv`, and a `PostRel` over recursive call events, `RPostInv`. Just like any form of invariants
698 in formal verification, correctly choosing `RPreInv` and `RPostInv` requires striking a careful
699 balance between choosing preconditions that are weak enough to hold, but strong enough to
700 imply post conditions. The rule is expressed in the following code.

```
701 Theorem padded_refines_mrec : forall (init1 : D1) (init2 : D2),
702   RPreInv init1 init2 →
703   (forall d1 d2, RPreInv d1 d2 →
704     padded_refines (SumRel RPreInv RPre)
705     (SumPostRel RPostInv RPost)
706     (RPostInv d1 d2)
707     (bodies1 d1) (bodies2 d2)) →
708   padded_refines RPre RPost (RPostInv init1 init2)
709   (mrec_spec bodies1 init1)
710   (mrec_spec bodies2 init2).
711
```

713 The hypotheses in this theorem state that the initial recursive calls, `init1` and `init2`, are in
714 the precondition `RPreInv`, and that given any two recursive calls related by the precondition,
715 `d1` and `d2`, the recursive function bodies refine one another, where recursive calls are related
716 by `RPreInv` and `RPostInv` and any other events are related by `RPre` and `RPost`. These reasoning
717 principles allow us to prove complicated propositions involving the coinductively defined
718 refinement relation without needing to perform direct coinduction.

719 While we include several parameter relations with the definition of `padded_refines`, at the
720 top level, we are typically interested in the case where all relations are set to equality. We
721 call this relation *strict refinement*, and refer to it with the \leq symbol.

```
722 Notation "phi1 '≤' phi2" :=
723   (padded_refines eq PostRelEq eq phi1 phi2).
724
```

726 Strict refinement is a transitive relation, and is strong enough to allow rewrites under the
727 context of any other application of `padded_refines`.

728 4.4 ITree specification Incompleteness

729 One way to interpret ITree specifications is as sets of ITrees. The following code defines
730 *concrete* ITree specifications, which correspond to executable ITrees.

```

731 Variant concreteF {E R} {EncodingType E} (F : itree_spec E R → Prop) : itree_spec E
732   R → Prop :=
733   | concreteRet (r : R) : concreteF F (Ret r)
734   | concreteTau (t : itree_spec E R) : F t → concreteF F (Tau t)
735   | concreteVis (e : E) (k : response_type e → itree_spec E R) :
736     (forall a, F (k a)) → concreteF F (Vis (Spec_vis e) k).
737 Definition concrete {E R} {EncodingType E} : itree_spec E R → Prop := gfp concreteF.
738
739

```

741 A concrete ITree specification contains no quantifiers along any of its branches. We can map
 742 each ITree specification to the set of `concrete` ITree specifications that refine it.

743 However, ITree specifications are not complete with respect to this interpretation. In
 744 particular, there are pairs of ITree specifications that represent equivalent sets of concrete
 745 ITree specifications, but do not refine one another. To see why, consider the following two
 746 ITree specification over an empty event signature `voidE`.

```

747 Definition top1 : itree_spec voidE unit :=
748   forall_spec void;; Ret tt.
749

```

```

751 Definition top2 : itree_spec voidE unit :=
752   or_spec spin (Ret tt).
753

```

755 Both `top1` and `top2` are refined by all concrete ITree specifications of type `itree_spec voidE unit`.
 756 We can prove the refinement for `top1` by applying the right `forall` rule, and reducing to a
 757 trivially satisfied proposition. For `top2`, we know that every concrete ITree specification of
 758 this type is `eutt` to either `spin` or `Ret tt`⁵. In each case, apply the right `exists` rule and
 759 choose the corresponding branch. However, given any relations `RE`, `REAns`, `RR`, we cannot
 760 prove `padded_refines RE REAns RR top1 top2`. This is because the only way to eliminate the
 761 `Spec_forall` on the left is to provide an element of the `void` type, which does not exist. This,
 762 along with the transitivity theorem, demonstrates that `padded_refines` is strictly weaker than
 763 the subset relation on sets of refining concrete ITree specification.

764 5 Total Correctness Specifications

765 This section discusses how to encode and prove simple pre- and post- condition specifications
 766 using ITree specifications. We also discuss how these definitions relate to our syntax-directed
 767 proof automation.

768 Suppose we have a program that takes in values of type `A` and returns values of type `B`.
 769 Suppose we want to prove that if given an input that satisfies a precondition `Pre : A → Prop`,
 770 it will return a value that satisfies a postcondition `Post : A → B → Prop` without triggering
 771 any other events. The postcondition is a relation over `A` and `B` to allow the postcondition to
 772 depend on the initial provided value. We can encode these conditions in the following ITree
 773 specification.

```

774 Definition total_spec : A → itree_spec E B :=
775   fun a => assume_spec (Pre a);;
776     b ← exists_spec B;;
777     assert_spec (Post a b);;
778     Ret b.
779
780

```

781 The specification assumes that the input satisfies the precondition, existentially introduces
 782 an output value, asserts the post condition holds, and finally returns the output.

⁵ Proving this fact requires a nonconstructive axiom like the Law of The Excluded Middle.

XX:20 Interaction Tree Specifications

```
Definition call_spec (a : A) : itree_spec (callE A B + E) B := trigger (inl (Call a)).

Definition calling' {F} `{EncodingType F} : (A → itree F B) →
  (forall (c : callE A B) , itree F (response_type c)) :=
  fun f c => f (unCall c).
Definition rec_spec (body : A → itree_spec (callE A B + E) B) (a : A) :
  itree_spec E B :=
  mrec_spec (calling' body) (Call a).
Definition rec_fix_spec
  (body : (A → itree_spec (callE A B + E) B) → A →
   itree_spec (callE A B + E) B) :
  A → itree_spec E B :=
  rec_spec (body call_spec).
```

■ Figure 11 `rec_fix_spec` Definition

783 The `total_spec` specification can be effectively used compositionally. Consider a merge
784 sort implementation, named `sort`, built on top of two recursively defined helper functions,
785 one for splitting a list in half, named `halve`, and one for merging sorted lists, named `merge`.
786 If we have already proven specializations of `total_spec` for these sub functions, it becomes
787 easier to prove a specification for `sort`. Immediately we can replace these sub functions with
788 their total correctness specification. Now consider how this total correctness specification
789 will behave on the left side of a refinement. First, we can eliminate `assume_spec (Pre a)` as
790 long as we can prove `Pre a`. Once we have done that, we get to universally introduce the
791 output `b`, along with a proof that it satisfies the post condition. We are finally left with only
792 `Ret b` with the assumption `Post a b`. This is a much simpler specification than our initial
793 executable specification, which relied on several control flow operators including a recursive
794 one.

795 However, this easy to use specification is not easy to directly prove. The `padded_refines_mrec`
796 rule gives us a sound reasoning principle for proving that a recursively defined function
797 refines another recursively defined function, but it does not give any direct insight into how
798 to prove any refinement that does not match that syntactic structure. To address this, we
799 introduce a recursively defined version of `total_spec_fix` that we can apply our recursive
800 reasoning principle on.

801 First, we introduce a specialization of the `mrec_spec` combinator called `rec_fix_spec`,
802 defined in Figure 11. The `rec_fix_spec` function has a type similar to that of a standard
803 fixpoint operator. The first argument, `body`, is a function that takes in a type of recursive
804 calls `A → itree_spec (callE A B + E) B` and an initial argument of type `A` and produces a
805 result in terms of an ITree specification. It relies on the `calling'` function to transform
806 this value into a value of type `forall (c:callE A B) , itree_spec (callE A B + E) B` which the
807 `mrec_spec` function requires. From there it relies on the `call_spec` and `rec_spec` functions to
808 wrap values of type `A` into `Call` events and `trigger` them.

809 Given this recursion operator, we introduce an equivalent version of the total correctness
810 specification, `total_spec_fix`.

```

811 Definition total_spec_fix : A → itree_spec E B :=
812   rec_fix_spec (fun rec a ⇒
813     assume_spec (Pre a);;
814     n ← exists_spec nat;;
815     trepeat n (
816       a' ← exists_spec A;;
817       assert_spec (Pre a' ∧ Rdec a' a);;
818       rec a'
819     );;
820     b ← exists_spec B;;
821     assert_spec (Post a b);;
822     Ret b).
823

```

825 This specification is reliant on the `trepeat n t` function, which simply binds an ITree, `t`, onto
826 the end of itself `n` times. Note that `total_spec_fix` is defined recursively, and contains the
827 elements of `total_spec` inside the recursive body. This makes it easier to relate to recursively
828 defined functions. It begins by assuming the precondition and ends by introducing an output,
829 asserting it satisfies the post condition, and returning the output. What comes between these
830 familiar parts requires more explanation. Recall the discussion of the `padded_refines_mrec`
831 rule. This reasoning principle lets you prove refinement between two recursively defined
832 ITree specifications when a single layer of unfolding of each specification match up one to
833 one with recursive calls.

834 This means that to have a useful, general, and recursively defined version of total
835 correctness specification we need to allow our recursive definition for total correctness
836 specification to choose the number of recursive calls the function requires. For this reason,
837 `total_spec_fix` existentially introduces a number `n` that specifies how many recursive calls are
838 needed for one level of unfolding of the recursive function starting at `a`. The specification then
839 includes `n` copies of a specification that existentially chooses a new argument `a'`, asserts a
840 predicate holds on it, and then recursively calls the specification on this new argument. This
841 asserted predicate contains two parts. First, we assert the precondition. A correct recursively
842 defined function should not call itself on an invalid input if given a valid input. Second, we
843 assert that `a'` is *less than* `a` according to the relation `Rdec`. In order for `total_spec_fix` to
844 actually be equivalent to `total_spec`, we need to assume that `Rdec` is well-founded⁶. The
845 fact that `Rdec` is well-founded ensures that this specification contains no infinite chains of
846 recursive calls. This allows us to prove that `total_spec_fix` refines `total_spec` as long as `Rdec`
847 is well-founded.

```

848 Theorem total_spec_fix_correct :
849   well_founded Rdec → forall (a : A), total_spec_fix a ≤ total_spec a.
850

```

852 This theorem allows us to initially prove refinement specifications for recursive functions
853 using the `padded_refines_mrec` rule with `total_spec_fix` and then replace it with the easier
854 to work with `total_spec`.

855 Both `total_spec` and `total_spec_fix` do not accept any ITree specifications that trigger
856 any events. As a result, these total correctness specifications do not allow any exceptions to
857 be raised, as you would expect with total correctness specifications.

858 5.1 Demonstration

859 To demonstrate how to work with `total_spec`, we describe how to verify the `merge` function,
860 a key component of the merge sort algorithm. The `merge` function takes two sorted lists

⁶ We use the Coq standard library's definition of well-foundedness for this.

XX:22 Interaction Tree Specifications

```
Definition merge : (list nat * list nat)
  →
  itree_spec E (list nat) :=
  rec_fix_spec (fun rec '(l1,l2) =>
    b1 ← is_nil l1;;
    b2 ← is_nil l2;;
    if b1 : bool then
      Ret l2
    else if b2 : bool then
      Ret l1
    else
      x ← head l1;;
      tx ← tail l1;;
      y ← head l2;;
      ty ← tail l2;;
      if Nat.leb x y then
        l ← rec (tx, y::ty);;
        Ret (x :: l)
      else
        l ← rec (x::tx, ty);;
        Ret (y::l)).

Definition merge_pre p :=
  let '(l1,l2) := p in
  sorted l1 ∧ sorted l2.
Definition merge_post '(l1,l2) l :=
  sorted l ∧ Permutation l (l1 ++ l2).

Definition rdec_merge '(l1,l2) '(l3,l4) :=
  length l1 < length l3 ∧
  length l2 = length l4 ∨
  length l1 = length l3 ∧
  length l2 < length l4.

Theorem merge_correct : forall l1 l2,
  merge (l1,l2) ≤ total_spec merge_pre
  merge_post (l1,l2).
```

■ **Figure 12** Merge implementation

861 and combines them into one larger sorted list which contains all the original elements. In
862 Figure 12, we present a recursively defined implementation of `merge` along with relevant
863 relations and the correctness theorem. The `merge` function is based on the standard list
864 manipulating functions `is_nil`, `head`, and `tail`. We assume that the event type `E` contains
865 some kind of error event which is emitted if `head` or `tail` is called on an empty list.⁷

866 The `merge` function relies on its arguments being sorted and guarantees that its output
867 is a single, sorted list that is a permutation of the concatenation of the original lists. We
868 formalize these conditions in `merge_pre` and `merge_post`. To prove that `merge` is correct, we
869 want to show that it refines the total specification built from its pre- and post- conditions.
870 To accomplish this, it suffices to choose a well founded relation and prove that `merge` satisfies
871 the resulting `total_spec_fix` specification. For this function, we use `rdec_merge` which ensures
872 that the pairs of lists that we recursively call `merge` on either both decrease in length, or one
873 decreases in length and the other has the same length.

874 This leaves us with a refinement goal between two recursively defined specifications. We
875 can then apply the `padded_refines_mrec_spec` theorem. For the relational precondition, we
876 require that each pair of `Call` events is equal, and that `Pre` holds on the value contained
877 within the call. For the relational postcondition, we require that equal `Call` events return
878 equal values and that `Post` holds on them. Finally, we can prove that the body `merge` refines
879 the body of `total_spec_fix` given these relation pre- and postconditions. We accomplish this
880 by setting the existential variables on the right to make a single recursive call and give it the
881 same argument as the recursive call that the body of `merge` makes.

882 With this technique, we can verify the simple server introduced in Section 1. Recall that
883 the `server_impl` program executes an infinite loop of receiving a list of numbers, sorting it,
884 and sending it back as a message. To verify `server_impl`, we first verify `halve`, the remaining
885 sub function of `sort`, using the same technique we used to prove the correctness of `merge`. We
886 can then use these facts to prove the correctness of `sort`, and use the correctness of `sort` to

⁷ We manage this assumption with a Coq type class called `ReSum`. For more information please read the original ITrees paper[30] or inspect the associated artifact.

Function Name	Description	C LoC	Proof LoC
<code>mbox_free_chain</code>	Deallocate an <code>mbox</code> chain	11	18
<code>mbox_len</code>	Compute the length in bytes of an <code>mbox</code> chain	9	40
<code>mbox_concat</code>	Concatenates an <code>mbox</code> chain after a single <code>mbox</code>	5	18
<code>mbox_concat_chains</code>	Concatenates two <code>mbox</code> chains	14	24
<code>mbox_split_at</code>	Split an <code>mbox</code> chain into two chains	25	147
<code>mbox_copy</code>	Copy a single <code>mbox</code>	13	74
<code>mbox_copy_chain</code>	Copy an <code>mbox</code> chain	18	173
<code>mbox_detach</code>	Detach the first <code>mbox</code> from a chain	18	18
<code>mbox_detach_from_end</code>	Detach the first N bytes from an <code>mbox</code> chain	3	50
<code>mbox_randomize</code>	Randomize the contents of an <code>mbox</code>	9	121
<code>mbox_drop</code>	Remove bytes from the start of an <code>mbox</code>	12	23

■ **Figure 13** Verified `mbox` functions

```

887 prove the correctness of server_impl.
888
889 Theorem server_correct :
890   (server_impl tt) ≤ (server_spec tt).
```

892 6 Automation and Evaluation

893 6.1 Auto-active Verification

894 A key goal of this work is to provide auto-active automation for ITree specifications refinement.
895 To this effect, the current section presents an automated Coq tactic for proving refinement
896 goals called `prove_refinement`. The `prove_refinement` tactic is designed to reduce proof goals
897 about refinement of programs to proof goals about the data and assertions used in those
898 programs. In the spirit of auto-active verification, this is done mostly automatically, but
899 with the user guiding the automation in places where human insight is needed.

900 The `prove_refinement` tactic defers to the user in two specific places. The first is in
901 defining invariants for uses of the `mrec` recursive function combinator. The tool defers to the
902 user to provide these invariants because inferring such invariants is undecidable. The second
903 place where `prove_refinement` defers to the user is in proving non-refinement goals regarding
904 first order data. The user can then apply other automated and/or manual proof techniques
905 for the theories of the resulting proof goals.

906 The `prove_refinement` tactic is defined using a collection of syntax-directed inference rules
907 for proving refinement goals. The tactic proves refinement goals by iteratively choosing and
908 applying a rule that matches the current goal and then proceeding to prove the antecedents.
909 The `prove_refinement` tactic implements this strategy using the Coq hint database mechanism,
910 which is already a user-extensible mechanism for proof automation using syntax-directed
911 rules.

912 We omit further implementation details both for space and because we do not claim the
913 implementation of the `prove_refinement` tactic is novel or interesting. What is novel and
914 interesting is that ITree specifications are designed in such a way that the straightforward
915 implementation is able to achieve impressive results.

916 **6.2 Evaluation**

917 He et al. [9] discussed using Heapster to verify the interface of `mbox`, a key datastructure in
 918 the implementation of the Encapsulating Security Payload (ESP) protocol of IPsec. The
 919 `mbox` datastructure represents a data packet as a linked list of fixed length arrays. He et al.
 920 [9] type checked and extracted functional specifications for several functions that manipulate
 921 `mbox`. Using ITree specifications, we specified and verified the behavior of these functional
 922 specifications using our auto-active verification tool. These functions are nontrivial, combining
 923 loops, recursion, and pointer manipulations. We present the list of verified functions in
 924 Figure 13.

925 For each function, we include the function's name, a description of its behavior, the
 926 number of lines of C code in its definition, and the number of lines of Coq code required
 927 to verify it. Lines of code are, of course, a very coarse metric for judging the complexity of
 928 code and proofs. However, these metrics do demonstrate the viability of this verification
 929 approach, showing that the remaining proof burden after the automation is of a reasonable
 930 size. The primary advantage this approach has over others is that the system reduces the
 931 verification down to facts about first order data. In this case, the data is a variant of the
 932 `mbox` datastructure written in Coq.

933 **7 Related Work**

934 The most closely related work is the work on Dijkstra monads [1, 16, 27, 28]. Dijkstra
 935 monads are a framework for writing specifications over arbitrary monads. This framework is
 936 the basis for verifying programs with effects in F^* [26], a programming language specifically
 937 designed for verification. Dijkstra monads arise from the interaction of three structures,
 938 a *monad* M , a *specification monad* W , and an *effect observation* function `obs`. The monad
 939 M represents computations to be verified, while the specification monad W is a monad for
 940 writing specifications about those computations. The effect observation function `obs` is a
 941 monad homomorphism that embeds computations in M to the most precise specification in
 942 W that they satisfy. The specification monad is also equipped with a refinement relation
 943 that expresses when one specification implies or is contained in another. As an example,
 944 Dijkstra monads arose out of generalizing the notion of weakest precondition computations,
 945 by viewing the weakest precondition transformer of a computation as itself being a stateful
 946 computation from postconditions to preconditions. The mapping from a computation to its
 947 weakest precondition transformer is then a monad homomorphism from the computation
 948 monad to the weakest precondition monad.

949 ITree specifications in fact form a Dijkstra monad, where the type `itree_spec E R` acts
 950 as the specification monad and the corresponding ITree monad `itree E R` without logical
 951 quantifier events forms the computation monad. The effect observation homomorphism is then
 952 the natural embedding from the ITree type without quantifiers to the type with quantifiers.
 953 Most Dijkstra monads are specialized to act as either partial specification logics, which
 954 always accept any nonterminating computations, or total specification logics, which always
 955 reject any nonterminating computations. This means that most existing Dijkstra monads
 956 cannot reason about termination-sensitive properties like liveness. ITree specifications have
 957 the advantage of admitting specifications that accept particular divergent computations and
 958 not others. For example, an ITree specification could accept any computation that produces
 959 an infinite pattern of messages and responses from a server, and reject any computation that
 960 silently diverges.

961 A notable exception is the work of Silver and Zdancewic [25], who also provided a Dijkstra

962 monad for ITrees. Much like ITree specifications it was capable of expressing specifications
963 that allow for specifying infinite behavior. However, it did not provide reasoning principles
964 for general recursion. The fact that ITree specifications represent specifications as syntax
965 rather than semantics, as an ITree rather than some function relating ITrees to `Prop`, enabled
966 us to write reasoning principles for general recursion and to build automation around the
967 refinement rules.

968 A lot of work on verifying monadic computations has been based on notions of equational
969 reasoning. This was in fact a key part of Moggi's original work [19]. Pitts [21] and Moggi [20]
970 extend this approach by building general theories of an evaluation predicate for reasoning
971 about return values of computations. This approach provides no explicit means to reason
972 about the effects, however, and also has no direct way of handling non-termination in
973 specifications such as the specifications needed for a server process. Plotkin and Pretnar [22]
974 further extend this approach with a general-purpose logic for algebraic effects, allowing it to
975 reason about the effects themselves and not just return values. This approach cannot handle
976 general Hoare logic assertions, however, and although there is a high-level discussion about
977 handling recursion, it is not clear how well it works for those sorts of specifications. Rauch
978 et al. [23] extends monads with native exceptions and non-termination and provides a logic
979 for these monads. Much like in our work, monads in Rauch et al. [23] can be annotated with
980 assertions. However, it restricts the language of assertions, and does not provide assumptions,
981 or general universal or existential quantification. It also handles only tail recursive programs,
982 and not general, mutual recursion.

983 One particularly effective approach in the space of equational reasoning was that of
984 Gibbons and Hinze [8]. This work showed how to use the specialized monad laws of each
985 sort of effect in a computation to define rewrite rules for simplifying and reasoning about
986 effectful computations, and then demonstrated that this approach is both straightforward to
987 use and powerful enough to verify a number of small but interesting programs.

988 The ultimate goal of this work is to provide techniques for auto-active verification of
989 imperative code. Therefore, it is natural to compare this work to semi-automated separation
990 logic tools like VST-Floyd[2] and CFML[6]. We argue this approach has two major advantages
991 over these related techniques. First, while VST-Floyd is specialized to C and CFML is
992 specialized to Caml, ITree specifications can be used to specify any programs with an
993 ITrees based semantics. When paired with Heapster techniques, ITree specifications can be
994 used to specify a wide array of imperative, heap-manipulating languages with a memory-safe
995 type system. In particular, the Heapster type system is closely related to the Rust type
996 system, meaning these techniques should be adaptable to specify and verify Rust code.
997 Second, the Heapster types are able to perform all the separation logic specific reasoning,
998 freeing the verifier to focus on the underlying mathematical structures.

999 **8 Conclusion**

1000 This paper introduces ITree specifications along with verified metatheory and proof automa-
1001 tion for reasoning about them. ITree specifications are a specialization of ITrees with a
1002 general notion of specification refinement. Unlike previous work developing specifications
1003 for ITrees, this paper provides techniques for working with the general recursion operator
1004 provided by the ITrees library. Finally, this paper demonstrates the effectiveness of its
1005 techniques by applying them on a collection of realistic C functions.

1006 References

- 1007 1 Danel Ahman, Catalin Hritcu, Kenji Maillard, Guido Martinez, Gordon Plotkin, Jonathan
1008 Protzenko, Aseem Rastogi, and Nikhil Swamy. Dijkstra monads for free. In *Proceedings of*
1009 *the 44th ACM SIGPLAN Symposium on Principles of Programming Languages (POPL)*,
1010 2017.
- 1011 2 Andrew W. Appel, Robert Dockins, Aquinas Hobor, Lennart Beringer, Josiah Dodds, Gor-
1012 don Stewart, Sandrine Blazy, and Xavier Leroy. *Program Logics for Certified Compilers*.
1013 Cambridge University Press, USA, 2014. ISBN 110704801X.
- 1014 3 Vytautas Astrauskas, Peter Müller, Federico Poli, and Alexander J. Summers. Leveraging
1015 rust types for modular specification and verification. In *Proceedings of the ACM SIGPLAN*
1016 *International Conference on Object-Oriented Programming, Systems, Languages, and*
1017 *Applications (OOPSLA)*, 2019.
- 1018 4 Armin Biere, Alessandro Cimatti, Edmund Clarke, and Yunshan Zhu. Symbolic model
1019 checking without bdds. In *Proceedings of the 5th International Conference on Tools and*
1020 *Algorithms for the Construction and Analysis of Systems (TACAS)*, 1999.
- 1021 5 Aaron R. Bradley. Sat-based model checking without unrolling. In *Proceedings of the 12th*
1022 *International Conference on Verification, Model Checking, and Abstract Interpretation*
1023 *(VMCAI)*, 2011.
- 1024 6 Arthur Charguéraud. Characteristic formulae for the verification of imperative programs.
1025 In *Proceedings of the 16th ACM SIGPLAN International Conference on Functional*
1026 *Programming, ICFP '11*, page 418–430, New York, NY, USA, 2011. Association for
1027 Computing Machinery. ISBN 9781450308656. doi: 10.1145/2034773.2034828. URL
1028 <https://doi.org/10.1145/2034773.2034828>.
- 1029 7 Andrey Chudnov, Nathan Collins, Byron Cook, Joey Dodds, Brian Huffman, Colm
1030 MacCárthaigh, Stephen Magill, Eric Mertens, Eric Mullen, Serdar Tasiran, Aaron Tomb,
1031 and Eddy Westbrook. Continuous formal verification of amazon s2n. In *Proceedings of*
1032 *the 30th International Conference on Computer Aided Verification (CAV)*, 2018.
- 1033 8 Jeremy Gibbons and Ralf Hinze. Just do it: simple monadic equational reasoning.
1034 In *Proceedings of the 16th ACM SIGPLAN international conference on Functional*
1035 *programming (ICFP)*, 2011.
- 1036 9 Paul He, Edwin Westbrook, Brent Carmer, Chris Phifer, Valentin Robert, Karl Smeltzer,
1037 Andrei Stefanescu, Aaron Tomb, Adam Wick, Matthew Yacavone, and Steve Zdancewic.
1038 A type system for extracting functional specifications from memory-safe imperative
1039 programs. In *Proceedings of the ACM SIGPLAN International Conference on Object-*
1040 *Oriented Programming, Systems, Languages, and Applications (OOPSLA)*, 2021.
- 1041 10 Son Ho and Jonathan Protzenko. Aeneas: Rust verification by functional translation.
1042 *Proc. ACM Program. Lang.*, 6(ICFP), aug 2022. doi: 10.1145/3547647. URL <https://doi.org/10.1145/3547647>.
- 1044 11 Chung-Kil Hur, Georg Neis, Derek Dreyer, and Viktor Vafeiadis. The power of
1045 parameterization in coinductive proof. In *Proceedings of the 40th Annual ACM*
1046 *SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, 2013. doi:
1047 10.1145/2429069.2429093.
- 1048 12 Ralf Jung, Robbert Krebbers, Lars Birkedal, and Derek Dreyer. Higher-order ghost state.
1049 In Jacques Garrigue, Gabriele Keller, and Eijiro Sumii, editors, *Proceedings of the 21st*
1050 *ACM SIGPLAN International Conference on Functional Programming, ICFP 2016, Nara,*
1051 *Japan, September 18-22, 2016*, pages 256–269. ACM, 2016. ISBN 978-1-4503-4219-3. doi:
1052 10.1145/2951913.2951943. URL <http://doi.acm.org/10.1145/2951913.2951943>.

- 1053 **13** Gerwin Klein, Kevin Elphinstone, Gernot Heiser, June Andronick, David Cock, Philip
1054 Derrin, Dhammika Elkaduwe, Kai Engelhardt, Rafal Kolanski, Michael Norrish, Thomas
1055 Sewell, Harvey Tuch, and Simon Winwood. sel4: Formal verification of an os kernel. In
1056 *Proceedings of the ACM SIGOPS 22Nd Symposium on Operating Systems Principles*,
1057 SOSP '09, pages 207–220, New York, NY, USA, 2009. ACM. ISBN 978-1-60558-752-3.
1058 doi: 10.1145/1629575.1629596. URL <http://doi.acm.org/10.1145/1629575.1629596>.
- 1059 **14** Xavier Leroy and Sandrine Blazy. Formal verification of a c-like memory model and its
1060 uses for verifying program transformations. *J. Autom. Reason.*, 41(1):1–31, jul 2008.
1061 ISSN 0168-7433. doi: 10.1007/s10817-008-9099-0. URL <https://doi.org/10.1007/s10817-008-9099-0>.
- 1063 **15** Giuliano Losa and Mike Dodds. On the Formal Verification of the Stellar Consensus
1064 Protocol. In *2nd Workshop on Formal Methods for Blockchains (FMBC 2020)*, 2020.
- 1065 **16** Kenji Maillard, Danel Ahman, Robert Atkey, Guido Martínez, Cătălin Hrițcu, Exequiel
1066 Rivas, and Éric Tanter. Dijkstra monads for all. *Proc. ACM Program. Lang.*, 3(ICFP),
1067 July 2019. doi: 10.1145/3341708. URL <https://doi.org/10.1145/3341708>.
- 1068 **17** Yusuke Matsushita, Takeshi Tsukada, and Naoki Kobayashi. Rusthorn: Chc-based
1069 verification for rust programs. In *Proceedings of the 29th European Symposium on*
1070 *Programming (ESOP)*, 2020.
- 1071 **18** Yusuke Matsushita, Xavier Denis, Jacques-Henri Jourdan, and Derek Dreyer. Rusthorn-
1072 belt: A semantic foundation for functional verification of rust programs with unsafe code.
1073 In *Proceedings of the 43rd ACM SIGPLAN Conference on Programming Language Design*
1074 *and Implementation*, 2022.
- 1075 **19** Eugenio Moggi. Computational lambda-calculus and monads. In *Proceedings of the*
1076 *Fourth Annual Symposium on Logic in Computer Science (LICS)*, 1989.
- 1077 **20** Eugenio Moggi. A semantics for evaluation logic. *Fundamenta Informaticae*, 22(1), 1989.
- 1078 **21** Andrew M. Pitts. Evaluation logic. In *Proceedings of the IV Higher Order Workshop*,
1079 1990.
- 1080 **22** Gordon Plotkin and Matija Pretnar. A logic for algebraic effects. In *Proceedings of the*
1081 *23rd Annual IEEE Symposium on Logic in Computer Science (LICS)*, 2008.
- 1082 **23** Christoph Rauch, Sergey Goncharov, and Lutz Schröder. Generic hoare logic for order-
1083 enriched effects with exceptions. In Phillip James and Markus Roggenbach, editors,
1084 *Recent Trends in Algebraic Development Techniques*, pages 208–222, Cham, 2017. Springer
1085 International Publishing. ISBN 978-3-319-72044-9.
- 1086 **24** Ilya Sergey, Aleksandar Nanevski, and Anindya Banerjee. Mechanized verification of
1087 fine-grained concurrent programs. In *Proceedings of the 36th ACM SIGPLAN Conference*
1088 *on Programming Language Design and Implementation, PLDI '15*, page 77–87, New
1089 York, NY, USA, 2015. Association for Computing Machinery. ISBN 9781450334686. doi:
1090 10.1145/2737924.2737964. URL <https://doi.org/10.1145/2737924.2737964>.
- 1091 **25** Lucas Silver and Steve Zdancewic. Dijkstra monads forever: Termination-sensitive
1092 specifications for interaction trees. *Proc. ACM Program. Lang.*, 5(POPL), jan 2021. doi:
1093 10.1145/3434307. URL <https://doi.org/10.1145/3434307>.
- 1094 **26** Nikhil Swamy, Juan Chen, Cédric Fournet, Pierre-Yves Strub, Karthikeyan Bharga-
1095 van, and Jean Yang. Secure distributed programming with value-dependent types. In
1096 *Proceeding of the 16th ACM SIGPLAN international conference on Functional Program-*
1097 *ming, ICFP 2011, Tokyo, Japan, September 19-21, 2011*, pages 266–278, 2011. doi:
1098 10.1145/2034773.2034811. URL <http://doi.acm.org/10.1145/2034773.2034811>.
- 1099 **27** Nikhil Swamy, Joel Weinberger, Cole Schlesinger, Juan Chen, and Benjamin Livshits.
1100 Verifying higher-order programs with the dijkstra monad. In *ACM SIGPLAN Conference*

- 1101 *on Programming Language Design and Implementation, PLDI '13, Seattle, WA, USA,*
1102 *June 16-19, 2013*, pages 387–398, 2013. doi: 10.1145/2491956.2491978. URL <http://doi.acm.org/10.1145/2491956.2491978>.
- 1103 **28** Nikhil Swamy, Cătălin Hrițcu, Chantal Keller, Aseem Rastogi, Antoine Delignat-Lavaud,
1104 Simon Forest, Karthikeyan Bhargavan, Cédric Fournet, Pierre-Yves Strub, Markulf
1105 Kohlweiss, Jean-Karim Zinzindohoue, and Santiago Zanella-Béguelin. Dependent types
1106 and multi-monadic effects in F^* . In *Proceedings of the 43rd Annual ACM SIGPLAN-*
1107 *SIGACT Symposium on Principles of Programming Languages*, POPL '16, page 256–270,
1108 New York, NY, USA, 2016. Association for Computing Machinery. ISBN 9781450335492.
1109 doi: 10.1145/2837614.2837655. URL <https://doi.org/10.1145/2837614.2837655>.
- 1110 **29** Li-yao Xia, Yannick Zakowski, Paul He, Chung-Kil Hur, Gregory Malecha, Benjamin C.
1111 Pierce, and Steve Zdancewic. Interaction trees: Representing recursive and impure
1112 programs in coq. *Proc. ACM Program. Lang.*, 4(POPL), dec 2019. doi: 10.1145/3371119.
1113 URL <https://doi.org/10.1145/3371119>.
- 1114 **30** Li-yao Xia, Yannick Zakowski, Paul He, Chung-Kil Hur, Gregory Malecha, Benjamin C.
1115 Pierce, and Steve Zdancewic. Interaction trees: Representing recursive and impure
1116 programs in coq. *Proceedings of the ACM on Programming Languages*, 4(POPL), January
1117 2020. doi: 10.1145/3371119.
- 1118