

Focusing on Refinement Typing

DIMITRIOS J. ECONOMOU, Queen's University, Canada

NEEL KRISHNASWAMI, University of Cambridge, United Kingdom

JANA DUNFIELD, Queen's University, Canada

We present a logically principled foundation for systematizing, in a way that works with any computational effect and evaluation order, SMT constraint generation seen in refinement type systems for functional programming languages. By carefully combining a focalized variant of call-by-push-value, bidirectional typing, and our novel technique of value-determined indexes, our system generates solvable SMT constraints without existential (unification) variables. We design a polarized subtyping relation allowing us to prove our logically focused typing algorithm is sound, complete, and decidable. We prove type soundness of our declarative system with respect to an elementary domain-theoretic denotational semantics. Type soundness implies, relatively simply, the total correctness and logical consistency of our system. The relative ease with which we obtain both algorithmic and semantic results ultimately stems from the proof-theoretic technique of focalization.

CCS Concepts: • **Theory of computation** → **Type theory; Type structures; Denotational semantics; Program reasoning; Proof theory**; • **Software and its engineering** → **Automated static analysis**.

Additional Key Words and Phrases: refinement types, bidirectional typechecking, polarity, call-by-push-value

1 INTRODUCTION

True, “well-typed programs cannot ‘go wrong’” [Milner 1978], but only relative to a given semantics (if the type system is proven sound with respect to it). Unfortunately, well-typed programs go wrong, in many ways that matter, but about which a conventional type system cannot speak: divisions by zero, out-of-bounds array accesses, information leaks. To prove a type system rules out (at compile time) more run-time errors, its semantics must be refined. However, there is often not enough type structure with which to express such semantics statically. So, we must refine our types with more information that tends to be related to programs. Great care is needed, though, because incorporating too much information (such as nonterminating programs themselves, as may happen in a *dependent* type system, where program terms may appear in types) can spoil good properties of the type system, like type soundness or the decidability of type checking and inference.

Consider the inductive type List A of lists with entries of type A . Such a list is either nil ($[]$) or a term x of type A together with a tail list xs (that is, $x :: xs$). In a typed functional language like Haskell or OCaml, the programmer can define such a type by specifying its constructors:

```
data List A where
  [] : List A
  (::) : A → List A → List A
```

Authors' addresses: Dimitrios J. Economou, Queen's University, Goodwin Hall 557, Kingston, ON, K7L 3N6, Canada, d.economou@queensu.ca; Neel Krishnaswami, University of Cambridge, Computer Laboratory, William Gates Building, Cambridge, CB3 0FD, United Kingdom, nk480@cl.cam.ac.uk; Jana Dunfield, Queen's University, Goodwin Hall 557, Kingston, ON, K7L 3N6, Canada, jd169@queensu.ca.

2023. 0164-0925/2023/8-ART \$15.00

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

Suppose we define, by pattern matching, the function `get`, that takes a list `xs` and a natural number `y`, and returns the `y`th element of `xs` (where the first element is numbered zero):

```
get [] y = error "Out of bounds"
get (x :: xs) zero = x
get (x :: xs) (succ y) = get xs y
```

A conventional type system has no issue checking `get` against, say, the type `List A → Nat → A` (for any type `A`), but `get` is unsafe because it throws an out-of-bounds error when the input number is greater than or equal to the length of the input list. If it should be impossible for `get` to throw such an error, then `get` must have a type where the input number is restricted to natural numbers strictly less than the length of the input list. Ideally, the programmer would simply refine the type of `get`, while leaving the program alone (except, perhaps, for omitting the first clause).

This, in contrast to dependent types [Martin-Löf 1984], is the chief aim of *refinement types* [Freeman and Pfenning 1991]: to increase the expressive power of a pre-existing (unrefined) type system, while keeping the latter’s good properties, like type soundness and, especially, decidability of typing, so that programmers are not too burdened with refactoring their code or manually providing tedious proofs. In other words, the point of refinement types is to increase the expressive power of a given type system while maintaining high automation (of typing for normal programs), whereas the point of dependent types is to be maximally expressive (even with the ambitious aim of expressing all mathematics), at the cost of automation (which dependent type system designers may try to increase after the fact of high expressivity).

To refine `get`’s type so as to rule out, statically, run-time out-of-bounds errors, we need to compare numbers against list lengths. Thus, we refine the type of lists by their length: $\{v : \text{List } A \mid \text{len } v = n\}$, the type of lists `v` of length `n`. This type looks a bit worrying, though, because the measurement, `len v = n`, seems to use a recursive program, `len`. The structurally recursive

```
len [] = 0
len (x :: xs) = 1 + len xs
```

happens to terminate when applied to lists, but there is no general algorithm for deciding whether an arbitrary computation terminates [Turing 1936]. As such, we would prefer not to use ordinary recursive programs at all in our type refinements. Indeed, doing so would seem to violate a phase distinction¹ [Moggi 1989a; Harper et al. 1990] between *static* (compile time) specification and *dynamic* (run time) program, which is almost indispensable for *decidable* typing.

The refinement type system Dependent ML (DML) [Xi 1998] provides a phase distinction in refining ML by an index domain which has no run-time content.² Type checking and inference in DML is only decidable when it generates constraints whose satisfiability is decidable. In practice, DML *did* generate decidable constraints, but that was not guaranteed by its design. DML’s distinction between indexes and programs allows it to support refinement types in the presence of computational effects (such as nontermination, exceptions, and mutable references) in a relatively straightforward manner. Further, the index-program distinction clarifies how to give a denotational semantics: a refinement type denotes a subset of what the type’s erasure (of indexes) denotes and a program denotes precisely what its erasure denotes [Melliès and Zeilberger 2015]. Dependent type systems, by contrast, do not have such an erasure semantics.

¹A language has a *phase distinction* if it can distinguish aspects that are relevant at run time from those that are relevant only at compile time.

²In today’s context, “Refinement ML” might seem a more appropriate name than Dependent ML. But when DML was invented, “refinement types” referred to *datasort* refinement systems; the abstract of Xi [1998] describes DML as “another attempt towards refining... type systems ..., following the step of refinement types (Freeman and Pfenning 1991).”

It seems liquid type systems [Rondon et al. 2008; Kawaguchi et al. 2009; Vazou et al. 2013, 2014] achieve highly expressive, yet sound and decidable *recursive refinements* [Kawaguchi et al. 2009] of inductive types by a kind of phase distinction: namely, by restricting the recursive predicates of specifications to terminating *measures* (like `len`) that soundly characterize, in a theory decidable by off-the-shelf tools like SMT solvers, the static structure of inductive types. Unlike DML, liquid typing can, for example, use the measure of whether a list of natural numbers is in increasing order, while remaining decidable. However, liquid typing’s lack of index-program distinction makes it unclear how to give it a denotational semantics, and has also led to subtleties involving the interaction between effects and evaluation strategy (we elaborate later in this section and Sec. 2). Vazou et al. [2014] appear to provide a denotational semantics in Section 3.3, but this is not really a denotational semantics in the sense we intend, because it is defined in terms of an operational semantics and not a separate and well-established mathematical model (such as domain theory).

Let’s return to the `get` example. Following the tradition of index refinement [Xi 1998], we maintain a phase distinction by syntactically distinguishing index terms, which can safely appear in types, from program terms, which cannot. In this approach, we want to check `get` against a more informative type

$$\forall l : \mathbb{N}. \underbrace{\{v : \text{List } A \mid \text{len } v = l\}}_{\text{List}(A)(l)} \rightarrow \{v : \text{Nat} \mid v < l\} \rightarrow A$$

quantifying over *indexes* l of sort \mathbb{N} (natural numbers) and requiring the accessing number to be less than l . However, this type isn’t quite right, because `Nat` is a type and \mathbb{N} is a sort, so writing “ $v < l$ ” confounds our phase distinction between programs and indexes. Instead, the type should look more like

$$\forall l : \mathbb{N}. \{v : \text{List } A \mid \text{len } v = l\} \rightarrow \{v : \text{Nat} \mid \text{index } v < l\} \rightarrow A$$

where

$$\begin{aligned} \text{index zero} &= 0 \\ \text{index (succ } y) &= 1 + \text{index } y \end{aligned}$$

computes the index term of sort \mathbb{N} that corresponds to a program term of type `Nat`, by a structural recursion homologous to that of `len`. The third clause of `get` has a nonempty list as input, so its index (length) must be $1 + m$ for some m ; the type checker assumes $\text{index (succ } y) < 1 + m$; by the aforementioned homology, these constraints are again satisfied at the recursive call ($\text{index } y < m$), until the second clause returns ($0 < 1 + m'$). The first clause of `get` is impossible, because no natural number is less than zero. We can therefore safely remove this clause, or (equivalently) replace `error` with `unreachable`, which checks against any type under a logically inconsistent context, such as $l : \mathbb{N}, n : \mathbb{N}, l = 0, n < l$ in this case.

$$\begin{aligned} \text{get } & \forall l, n : \mathbb{N}. \{v : \text{List } A \mid \text{len } v = l\} \rightarrow \{v : \text{Nat} \mid \text{index } v = n\} \wedge (n < l) \rightarrow A \\ \text{get } [] & y = \text{unreachable} \quad \text{-- } l = 0 \text{ and } n : \mathbb{N} \text{ so } n \not< l \\ \text{get } (x :: xs) & \text{zero} = x \\ \text{get } (x :: xs) & (\text{succ } y) = \text{get } xs \ y \end{aligned}$$

Applying `get` to a list and a number determines the indexes l and n . We say that l and n are *value-determined* (here by applying the function to values). If (perhaps in a recursive call) `get` is called with an empty list `[]` and a natural number, then l is determined to be 0, and since no index that is both negative and a natural number exists, no out-of-bounds error can arise by calling `get`. (Further, because $l : \mathbb{N}$ strictly decreases at recursive calls, calling `get` terminates.)

While this kind of reasoning about `get`’s type refinement may seem straightforward, how do we generalize it to recursion over any algebraic datatype (ADT)? What are its logical and semantic

ingredients? How do we use these ingredients to concoct a type system with decidable typing, good (localized) error messages and so on, while also keeping its metatheory relatively stable or maintainable under various extensions or different evaluation strategies?

Type systems that can do this kind of reasoning automatically, especially in a way that can handle any evaluation strategy, are hard to design correctly. Indeed, the techniques used in the original (call-by-value) liquid type system(s) [Rondon et al. 2008; Kawaguchi et al. 2009] had to be modified for Haskell, essentially because of Haskell’s *call-by-name* evaluation order [Vazou et al. 2014]. The basic issue was that binders can bind in (static) refinement predicates, which is fine when binders only bind values (as in call-by-value), but not when they bind computations which may not terminate (as in call-by-name). Liquid Haskell regained (operational) type soundness by introducing ad hoc restrictions that involve approximating whether binders terminate, and using the refinement logic to verify termination.

We design a foundation on which to build liquid typing features that allows us to establish clear (semantic) soundness results, as well as the completeness of a decidable bidirectional typing algorithm. The main technique facilitating this is focusing, which we combine with bidirectional typing and value-determined indexes (the latter being a key ingredient to make measures work). In other words, this paper is a first step toward reconciling DML and Liquid Haskell, using the proof-theoretic technique of focusing.

Andreoli [1992] introduced *focusing* to reduce the search space for proofs (programs) of logical formulas (types), by exploiting the property that some inference rules are invertible (the rule’s conclusion implies its premises). In relation to functional programming, focusing has been used, for example, to explain features such as pattern matching [Krishnaswami 2009], to explain the interaction between evaluation order and effects [Zeilberger 2009], and to reason about contextual program equivalence [Rioux and Zdancewic 2020]. Focusing has been applied to design a union and intersection refinement typing algorithm [Zeilberger 2009]. As far as we know, until now focusing has not been used to design an *index* refinement typing algorithm. By focusing on the typing of function argument lists and results, our focused system guarantees that value-determined existential indexes (unification variables) are solved before passing constraints to an SMT solver. For example, when our system infers a type for `get([3, 1, 2], 2)`, we first use the top-level annotation of `get` to synthesize the type

$$\downarrow (\forall l, n : \mathbb{N}. \{v : \text{List } A \mid \text{len } v = l\} \rightarrow \{v : \text{Nat} \mid \text{index } v = n\} \wedge (n < l) \rightarrow \uparrow A)$$

(in which we have added polarity shifts \downarrow – and \uparrow – arising from focusing). The downshift \downarrow – takes a negative type to a positive type of suspended computations. Second, we check the argument list `([3, 1, 2], 2)` against the negative (universally quantified) type. The upshift \uparrow – takes a positive type A to negative type $\uparrow A$ (computations returning a value of type A). In typechecking the argument list, the upshift signifies the end of a (left) focusing stage, at which point the first argument value `[3, 1, 2]` will have determined l to be 3 and the second argument value `2` will have determined n to be 2, outputting an SMT constraint without existentials: $2 < 3$.

Levy [2004] introduced the paradigm and calculus *call-by-push-value* (CBPV) which puts both call-by-name and call-by-value on equal footing in the storm of computational effects (such as nontermination). CBPV subsumes both call-by-name (CBN) and call-by-value (CBV) functional languages, because it allows us to encode both via type discipline. In particular, CBPV polarizes types into (positive) value types P and (negative) computation types N , and provides polarity shifts $\uparrow P$ (negative) and $\downarrow N$ (positive); the monads functional programmers use to manage effects arise as the composite $\downarrow \uparrow$ –. These polarity shifts are the same as those arising from focusing. CBPV can be derived logically by way of focalization [Espírito Santo 2017], which we did in our system. Focalized CBPV is a good foundation for a refinement typing algorithm: designing refinement

typing algorithms is challenging and sensitive to effects and evaluation strategy, so it helps to refine a language that makes evaluation order explicit. We leverage focusing and our technique of value-determined indexes (a novelty in the DML tradition) to guarantee (like Liquid Haskell) the generation of SMT-solvable constraints.

Bidirectional typing [Pierce and Turner 2000] systematizes the difference between input (for example, type checking) and output (for example, type inference), and seems to fit nicely with focused systems [Dunfield and Krishnaswami 2021]. Bidirectional typing has its own practical virtues: it is easy to implement (if inputs and outputs fit together properly, that is, if the system is well-moded); it scales well (to refinement types, higher-rank polymorphism [Dunfield and Krishnaswami 2019], subtyping, effects—and so does CBPV); it leads to localized error messages; and it clarifies where type annotations are needed, typically in reasonable positions (such as at the top level) that are helpful as machine-checked documentation. In our system, annotations are needed only for recursive functions (to express termination refinements) and top-level definitions.

A focused and bidirectional approach therefore appears suitable, both theoretically and practically, for systematically designing and implementing an expressive language of type refinement that can handle any evaluation strategy and effect. We show that bidirectional typing and logical focusing work very well together at managing the complex flow of information pertaining to indexes of recursive data. In particular, value-determined existential indexes of input types are solved within focusing stages, ultimately leading to the output of constraints (and types) in the quantifier-free fragment solvable by SMT solvers.

Contributions. Our two key contributions are *both* a declarative/logical/semantic *and* an algorithmic account of recursive, index-based refinement of algebraic data types. For the logical account, we design a declarative type system in a *bidirectional* and *focused* style, resulting in a system with clear denotational semantics and soundness proofs, and which is convenient for type theorists of programming languages. The declarative system conjures index solutions to existentials. For the algorithmic account, we design a type system similar to the declarative one but solving all existentials, and prove it is decidable, as well as sound and complete. We contribute:

- A polarized declarative type system, including (polarized) subtyping, universal types, existential types, and index refinements with ordinary SMT constraints, as well as (nullary) recursive predicates on inductive data (which in ongoing work we are extending to multi-argument measures, which can express, for example, lists of increasing integer elements).
- A proof that declarative typing is stable under substitution, which requires proving, among other things, that subtyping is transitive and typing subsumption is admissible.
- A clear denotational semantics of the declarative system, based on elementary domain theory.
- A type soundness proof with respect to our denotational semantics, which implies, relatively easily, both the refinement system’s logical consistency and total correctness—even if the programs are *non-structurally* recursive. To prove type soundness, we prove that value-determined indexes are sound: that is, semantic values uniquely determine value-determined indexes, semantically speaking (in particular, see Lemma 5.4).
- A polarized subtyping algorithm, together with proofs that it is sound, complete and decidable.
- A polarized typing algorithm, together with proofs that it is sound, complete and decidable. Completeness relies on the combination of our novel technique of *value-determined* indexes, focusing, and bidirectional typing. In particular, Lemma 7.7 implies that all existential variables are solved by the algorithm.

We relatively easily obtain both semantic and algorithmic results for a realistic language essentially by applying just one technique (based on fundamental logical principles): focusing.

All proofs are in the appendix.

2 OVERVIEW

This paper is a first step toward reconciling Dependent ML and Liquid Haskell. The main thing we get from DML is the index-program distinction. Liquid Haskell provides or inspires three things. First, the observation of difficulties with effects and evaluation order inspired our use of CBPV. Second, we study (nullary) measures (supporting multi-argument measures is ongoing work). Third, our technique of value-determined indexes was inspired by the observation that variables appearing in liquid refinements correspond to inputs or outputs of functions.

Before diving into the details of our type system, we give an overview of the central logical, semantic, and algorithmic issues informing its design. The main technique we use to easily support both semantic and algorithmic results is focalization.

Refinement typing, evaluation strategy, and computational effects. The interactions between refinement typing (and other fancy typing), evaluation strategy, and computational effects are a source of peril. The combination of parametric polymorphism with effects is often unsound [Harper and Lillibridge 1991]; the value restriction in Standard ML recovers soundness in the presence of mutable references by restricting polymorphic generalization to syntactic values [Wright 1995]. The issue was also not peculiar to polymorphism: Davies and Pfenning [2000] discovered that a similar value restriction recovers type soundness for intersection refinement types and effects in call-by-value languages. For union types, Dunfield and Pfenning [2003] obtained soundness by an evaluation context restriction on union elimination.

For similar reasons, Liquid Haskell was also found unsound in practice, and had to be patched; we adapt an example [Vazou et al. 2014] demonstrating the discovered unsoundness:

```

diverge :: Nat -> {v:Int | false}
diverge x = diverge x

safediv :: n:Nat -> {d:Nat | 0 < d} -> {v:Nat | v <= n}
safediv n d = if 0 < d then n / d else error "unreachable"

unsafe :: Nat -> Int
unsafe x = let notused = diverge 1 in let y = 0 in safediv x y

```

In typechecking `unsafe`, we need to check that the type of `y` (a singleton type of one value: 0) is a subtype of `safediv`'s second argument type (under the context of the `let`-binding). Due to the refinement of the `let`-bound `notused`, this subtyping generates a *constraint* or *verification condition* of the form “if `false` is true, then...”. This constraint holds vacuously, implying that `unsafe` is safe. But `unsafe` really is unsafe because Haskell evaluates *lazily*: since `notused` is not used, `diverge` is never called, and hence `safediv` divides by zero (and crashes if uncaught). Vazou et al. [2014] recover type soundness and decidable typing by restricting `let`-binding and subtyping, using an operational semantics to approximate whether or not expressions diverge, and whether or not terminating terms terminate to a *finite* value.

The value and evaluation context restrictions seem like ad hoc ways to cope with the failure of simple typing rules to deal with the interactions between effects and evaluation strategy. However, Zeilberger [2009] explains the value and evaluation context restrictions in terms of a logical view of refinement typing. Not only does this perspective explain these restrictions, it provides theoretical tools for designing type systems for functional languages with effects. At the heart of Zeilberger's approach is the proof-theoretic technique of *focusing*, which we discuss near the end of this overview. An important question we address is whether polarization and focusing can also help us understand Liquid Haskell's restrictions on `let`-binding and subtyping: basically, our `let`-binding rule requires

the bound computation (negative type) to terminate to a value (positive type). In other words, focalized systems satisfy any necessary value (and covalue) restrictions by default. We discuss this further in Sec. 8.

Focalization can also yield systems with good semantic properties under computational effects, in particular, variants of call-by-push-value.

Refining call-by-push-value. Call-by-push-value [Levy 2004] subsumes both call-by-value and call-by-name by polarizing the usual terms and types of the λ -calculus into a finer structure that can be used to encode both evaluation strategies in a way that can accommodate computational effects: *value* (or *positive*) types (classifying terms which “are”, that is, values v), *computation* (or *negative*) types (classifying terms which “do”, that is, expressions e), and polarity shifts $\uparrow-$ and $\downarrow-$ between them. An upshift $\uparrow P$ lifts a (positive) value type P up to a (negative) computation type of expressions that compute values (of type P). A downshift $\downarrow N$ pushes a (negative) computation type N down into a (positive) value type of *thunked* or *suspended* computations (of type N). We can embed the usual λ -calculus function type $A \rightarrow_{\lambda} B$ (written with a subscript to distinguish it from the CBPV function type), for example, into CBPV (whose function types have the form $P \rightarrow N$ for positive P and negative N) so that it behaves like CBV, via the translation ι_{CBV} with $\iota_{\text{CBV}}(A \rightarrow_{\lambda} B) = \downarrow(\iota_{\text{CBV}}(A) \rightarrow \uparrow\iota_{\text{CBV}}(B))$; or so that it behaves like CBN, via the translation ι_{CBN} with $\iota_{\text{CBN}}(A \rightarrow_{\lambda} B) = (\downarrow\iota_{\text{CBN}}(A)) \rightarrow \iota_{\text{CBN}}(B)$.

Evaluation order is made explicit by CBPV type discipline. Therefore, adding a refinement layer on top of CBPV requires directly and systematically dealing with the interaction between type refinement and evaluation order. If we add this layer to CBPV correctly from the very beginning, then we can be confident that our type refinement system will be semantically well-behaved when extended with other computational effects. The semantics of CBPV are well-studied and this helps us establish semantic metatheory. In later parts of this overview, we show the practical effects of refining our focalized variant of CBPV, especially when it comes to algorithmic matters.

Type soundness, totality, and logical consistency. The unrefined system underlying our system has the computational effect of nontermination and hence is not total. To model nontermination, we give the unrefined system an elementary domain-theoretic denotational semantics. Semantic type soundness says that a syntactic typing derivation can be faithfully interpreted as a semantic typing derivation, that is, a morphism in a mathematical category, in this case a logical refinement of domains. Semantic type soundness basically corresponds to syntactic type soundness with respect to a big-step operational semantics. While we don’t provide an operational semantics in this paper, we do prove a syntactic substitution lemma which would be a key ingredient to prove that an operational semantics preserves typing (beta reduction performs syntactic substitution). The substitution lemma is also helpful to programmers because it means they can safely perform program transformations and preserve typing. Because the unrefined system is (a focalized variant of) CBPV, proving type soundness is relatively straightforward.

In contrast to dependent types, the denotational semantics of our refined system is defined in terms of that of its *erasure* (of indexes), that is, its underlying, unrefined system. A refined type denotes a logical subset of what its erasure denotes. An *unrefined* return type $\uparrow P$ denotes either what P denotes, or divergence/nontermination. A *refined* return type $\uparrow P$ denotes *only* what P denotes. Therefore, our *refined* type soundness result implies that our refined system (without a *partial* upshift type) enforces termination. In we discuss how to extend the refined system (by a *partial* upshift type) to permit divergence while keeping type soundness (which implies partial correctness for *partial* upshifts). Type soundness also implies logical consistency, because a logically inconsistent refinement type denotes the empty set. We also prove that syntactic substitution is

semantically sound, which would be a main lemma in proving that an operational semantics is equivalent to our denotational semantics.

In Sec. 5, we discuss these semantic issues in more detail.

Algebraic data types and measures. A novelty of liquid typing is the use of *measures*: functions, defined on algebraic data types, which may be structurally recursive, but are guaranteed to terminate and can therefore safely be used to refine the inductive types over which they are defined. (In this paper, we only consider nullary measures.)

For example, consider the type `BinTree A` of binary trees with terms of type `A` at leaves:

```
data BinTree A where
  leaf : A → BinTree A
  node : BinTree A → BinTree A → BinTree A
```

Suppose we want to refine `BinTree A` by the height of trees. Perhaps the most direct way to specify this is to measure it using a function `hgt` defined by structural recursion:

```
hgt : BinTree A → ℕ
hgt leaf = 0
hgt (node t u) = 1 + max(hgt(t), hgt(u))
```

As another example, consider an inductive type `Expr` of expressions in a CBV lambda calculus:

```
data Expr where
  var : Nat → Expr
  lam : Nat → Expr → Expr
  app : Expr → Expr → Expr
```

Measures need not involve recursion. For example, if we want to refine the type `Expr` to expressions `Expr` that are values (in the sense of CBV, not CBPV), then we can use `isval`:

```
isval : Expr → ℬ
isval (var z) = tt
isval (lam z expr) = tt
isval (app expr expr') = ff
```

Because `isval` isn't recursive and returns indexes, it's safe to use it to refine `Expr` to expressions that are CBV values: $\{v : \text{Expr} \mid \text{isval } v = \text{tt}\}$. But, as with `len` (Sec. 1), we may again be worried about using the seemingly dynamic, recursively defined `hgt` in a static type refinement. Again, we need not worry because `hgt`, like `len`, is a terminating function into a decidable logic [Barrett et al. 2009]. We can use it to specify that, say, a height function defined by pattern matching on trees of type $\{v : \text{BinTree } A \mid \text{hgt } v = n\}$ actually returns (the program value representing) n for any tree of height n . Given the phase distinction between indexes (like n) and programs, how do we specify such a function type? We use refinement type unrolling and singletons.

Unrolling and singletons. Let's consider a slightly simpler function, `length`, that takes a list and returns its length:

```
length [] = zero
length (x :: xs) = succ (length xs)
```

What should be the type specifying that `length` actually returns a list's length? The proposal $\forall n : \mathbb{N}. \text{List}(A)(n) \rightarrow \uparrow \text{Nat}$ does not work because `Nat` has no information about the length n . Something like $\forall n : \mathbb{N}. \text{List}(A)(n) \rightarrow \uparrow (n : \text{Nat})$, read literally as returning the index n , would violate our phase distinction between programs and indexes. Instead, we use a *singleton* type in the sense of Xi [1998]: a singleton type contains just those program terms (of the type's erasure), that

correspond to exactly one semantic index. For example, given $n : \mathbb{N}$, we define the singleton type $\text{Nat}(n)$ (which may also be written $\text{Nat } n$) by $\{v : \text{Nat} \mid \text{index } v = n\}$ where

$$\begin{aligned} \text{index} &: \text{Nat} \rightarrow \mathbb{N} \\ \text{index zero} &= 0 \\ \text{index (succ } x) &= 1 + \text{index}(x) \end{aligned}$$

specifies the indexes (of sort \mathbb{N}) corresponding to program values of type Nat .

How do we check length against $\forall n : \mathbb{N}. \text{List}(A)(n) \rightarrow \uparrow(\text{Nat}(n))$? In the first clause, the input $[]$ has type $\text{List}(A)(n)$ for some n , but we need to know $n = 0$ (and that the index of zero is 0). Similarly, we need to know $x :: xs$ has length $n = 1 + n'$ where $n' : \mathbb{N}$ is the length of xs . To generate these constraints, we use an *unrolling* judgment (Sec. 4.6) that unrolls a refined inductive type. Unrolling puts the type's refinement constraints, expressed by *asserting* and *existential* types, in the structurally appropriate positions. An asserting type is written $P \wedge \phi$ (read “ P with ϕ ”), where P is a (positive) type and ϕ is an index proposition. If a term has type $P \wedge \phi$, then the term has type P and also ϕ holds. (Dual to asserting types, we have the *guarded* type $\phi \supset N$, which is equivalent to N if ϕ holds, but is otherwise useless.) We use asserting types to express that index equalities like $n = 0$ hold for terms of inductive type. We use existentials to quantify over indexes that characterize the refinements of recursive subparts of inductive types, like n' . For example, modulo a small difference (see Sec. 4.6), $\text{List}(A)(n)$ unrolls to

$$(1 \wedge (n = 0)) + (A \times \exists n' : \mathbb{N}. \{v : \text{List } A \mid \text{len } v = n'\} \wedge (n = 1 + n'))$$

That is, to construct an A -list of length n , the programmer (or library designer) can either left-inject a unit value, *provided the constraint $n = 0$ holds*, or right-inject a pair of one A value and a tail list, *provided that n' , the length of the tail list, is $n - 1$* (the equations $n = 1 + n'$ and $n - 1 = n'$ are equivalent). These index constraints are not a syntactic part of the list itself. That is, a term of the above refined type is also a term of the type's *erasure* (of indexes):

$$1 + (|A| \times (\text{List } |A|))$$

where $| - |$ erases indexes. Dual to *verifying* refined inductive values, pattern matching on refined inductive values, such as in the definition of length, allows us to *use* the index refinements locally in the bodies of the various clauses for different patterns. Liquid Haskell similarly extracts refinements of data constructors for use in pattern matching.

The shape of the refinement types generated by our unrolling judgment (such as the one above) is a judgmental and refined-ADT version of the fact that generalized ADTs (GADTs) can be desugared into types with equalities and existentials that express constraints of the return types for constructors [Cheney and Hinze 2003; Xi et al. 2003]. It would be tedious and inefficient for the programmer to work directly with terms of types produced by our unrolling judgment, but we can implement (in our system) singleton refinements of base types and common functions on them, such as addition, subtraction, multiplication, division, and the modulo operation on integers, and build these into the surface language used by the programmer, similarly to the implementation of Dependent ML [Xi 1998].

Inference and subtyping. For a typed functional language to be practical, it must support some degree of inference, especially for function application (to eliminate universal types) and constructing values (to introduce existential types). For example, to pass a value to a function, its type must be compatible with the function's argument type, but it would be burdensome to make programmers always have to prove this compatibility. In our setting, for example, if $x : \text{Nat}(3)$ and $f : \downarrow(\forall a : \mathbb{N}. \text{Nat}(a) \rightarrow \uparrow P)$, then we would prefer to write $f x$ rather than $f [3] x$, which would quickly make our programs incredibly—and unnecessarily—verbose.

Omitting index and type annotations, however, has significant implications. In particular, we need a mechanism to instantiate indexes somewhere in our typing rules: for example, if $g : \downarrow(\downarrow(\text{Nat}(4 + b) \rightarrow \uparrow P) \rightarrow N)$ and $h : \downarrow((\exists a : \mathbb{N}. \text{Nat}(a)) \rightarrow \uparrow P)$, then to apply g to h , we need to know $\text{Nat}(4 + b)$ is compatible with $\exists a : \mathbb{N}. \text{Nat}(a)$, which requires instantiating the bound a to a term logically equal to $4 + b$. Our system does this kind of instantiation via subtyping, which refinement types naturally give rise to: a type refinement is essentially a subtype of its erasure. Index instantiations are propagated locally across adjacent nodes in the syntax tree, similarly to Pierce and Turner [2000]. (Liquid typing allows for more inference, including inference of refinements based on templates, which provides additional convenience for programmers, but we do not consider this kind of inference in this paper.)

We polarize subtyping into two, mutually recursive, positive and negative relations $\Theta \vdash P \leq^+ Q$ and $\Theta \vdash N \leq^- M$ (where Θ is a *logical context* including index propositions). The algorithmic versions of these only introduce existential variables in positive supertypes and negative subtypes, guaranteeing they can always be solved by indexes without any existential variables. We delay checking constraints until the end of certain, logically designed stages (the *focusing* ones, as we will see), when all of their existential variables are guaranteed to have been solved.

Value-determined indexes and type well-formedness. Like DML, we have an index-program distinction, but unlike DML and like Liquid Haskell, we want to guarantee SMT solvable constraints. We accomplish this with our technique of value-determined indexes. To guarantee that our algorithm can always instantiate quantifiers, we restrict quantification to indexes appearing in certain positions within types: namely, those that are uniquely determined (semantically speaking) by values of the type, both according to a measure and before crossing a polarity shift (which in this case marks the end of a focusing stage). For example, in $\{\nu : \text{List } A \mid \text{len } \nu = b\}$, the index b is uniquely determined by values of that type: the list $[x, y]$ uniquely determines b to be 2 (by the length measure). This value-determinedness restriction on quantification has served to explain why a similar restriction in the typing algorithm of Flux (Liquid Rust) seemed to work well in practice [Lehmann et al. 2023].

We make this restriction in the type well-formedness judgment, which outputs a context Ξ tracking value-determined indexes; well-formed types can only quantify over indexes in Ξ . For example, $\exists b : \mathbb{N}. \{\nu : \text{List } A \mid \text{len } \nu = b\}$ is well-formed. The variables in Ξ also pay attention to type structure: for example, a value of a product type is a pair of values, where the first value determines all Ξ_1 (for the first type component) and the second value determines all Ξ_2 (second type component), so the Ξ of the product type is their union $\Xi_1 \cup \Xi_2$. We also take the union for function types $P \rightarrow N$, because index information flows through argument types toward the return type, marked by a polarity shift.

By emptying Ξ at shift types, we prevent lone existential variables from being introduced at a distance, across polarity shifts. In practice, this restriction on quantification is not onerous, because most functional types that programmers use are, in essence, of the form

$$\forall \dots. P_1 \rightarrow \dots \rightarrow P_n \rightarrow \uparrow \exists \dots. Q$$

where the “ \forall ” quantifies over indexes of argument types P_k that are uniquely determined by argument values, and the “ \exists ” quantifies over indexes of the return type that are determined by (or at least depend on) fully applying the function and thereby constructing a value to return. The idea of this restriction was inspired by liquid types because they implicitly follow it: variables appearing in liquid type refinements must ultimately come from arguments x to dependent functions $x : A \rightarrow B$ and their return values (however, these are not explicitly index variables).

Types that quantify only across polarity shifts tend to be empty, useless, or redundant. The ill-formed type $\forall n : \mathbb{N}. 1 \rightarrow \uparrow \text{Nat}(n)$ is empty because no function returns all naturals when applied to unit. A term of ill-formed type $\exists m : \mathbb{N}. \downarrow (\text{Nat}(m) \rightarrow \uparrow \text{Bool})$ can only be applied to an unknown number, which is useless because the number is unknown. The ill-formed type $\exists n : \mathbb{N}. \uparrow \downarrow \text{Nat}(n)$ is redundant because it is semantically equivalent to $\downarrow \uparrow \exists n : \mathbb{N}. \text{Nat}(n)$ (which does not quantify across a polarity shift), and similarly $\forall n : \mathbb{N}. \uparrow \downarrow (\text{Nat}(n) \rightarrow \uparrow \text{Nat}(n))$ is semantically equivalent to $\uparrow \downarrow (\forall n : \mathbb{N}. \text{Nat}(n) \rightarrow \uparrow \text{Nat}(n))$. Some refinements are *not* value-determined but useful nonetheless, such as dimension types [Kennedy 1994; Dunfield 2007b] which statically check that dimensions are used consistently (minutes can be added to minutes, but not to kilograms) but do not store the dimensions at run time. In this paper, we do not consider these non-value-determined refinements, and Liquid Haskell does not support them either.

Our value-determinedness restriction on type well-formedness, together with focusing, is very helpful metatheoretically, because it means that our typing algorithm only introduces—and is guaranteed to solve—existential variables for indexes within certain logical stages. For example, consider checking a list against $\exists b : \mathbb{N}. \{v : \text{List } A \mid \text{len } v = b\}$. An existential variable \hat{b} for b is generated, and we check the unrolled list against the unrolling of $\{v : \text{List } A \mid \text{len } v = \hat{b}\}$. A solution to \hat{b} will be found within value typechecking (the right-focusing stage), using the invariant that no measure (such as `len`) contains any existential variables. Similarly, applying a function with universal quantifiers will solve all existential variables arising from these quantifiers by the end of a left-focusing stage, which typechecks an argument list.

Focusing, CBPV, and bidirectional typing. In proof theory, the technique of *focusing* [Andreoli 1992] exploits invertibility properties of logical formulas (types), as they appear in inference rules (typing rules), in order to rule out many redundant proofs. Having fewer possible proofs makes proof search more tractable. Brock-Nannestad et al. [2015] and Espírito Santo [2017] study the relation between CBPV and focusing: each work provides a focused calculus that is essentially the same as CBPV, “up to the question of η -expansion” [Brock-Nannestad et al. 2015]. Our system is also a focused variant of CBPV; in fact, it arises from a certain focalization (and bidirectionalization) of ordinary intuitionistic logic.

An inference rule is *invertible* if its conclusion implies its premises. For example, in intuitionistic logic, the right rule for implication is invertible because its premise $\Gamma, A \vdash B$ can be derived from its conclusion $\Gamma \vdash A \rightarrow B$:

$$\frac{\frac{\frac{}{\Gamma \vdash A \rightarrow B} \text{ (Assume } \rightarrow\text{R conclusion)}}{\Gamma, A \vdash A \rightarrow B} \text{ (Weaken)}}{\Gamma, A \vdash B} \quad \frac{\frac{\frac{}{\Gamma, A \vdash A} \quad \frac{}{\Gamma, A, B \vdash B}}{\Gamma, A, A \rightarrow B \vdash B} \rightarrow\text{L}}{\Gamma, A \vdash B} \text{ (Cut)}$$

However, both right rules for disjunction, for example, are not invertible, which we can prove with a counterexample: $A_1 + A_2 \vdash A_1 + A_2$ but $A_1 + A_2 \not\vdash A_1$ and $A_1 + A_2 \not\vdash A_2$. In a sequent calculus, *positive* formulas have invertible left rules and *negative* formulas have invertible right rules. A *weakly* focused sequent calculus eagerly applies non-invertible rules as far as possible (in either *left-* or *right-focusing* stages); a *strongly* focused sequent calculus does too, but also eagerly applies invertible rules as far as possible (in either *left-* or *right-inversion* stages). There are also *stable* stages (or moments) in which a decision has to be made between focusing on the left, or on the right [Espírito Santo 2017]. The decision can be made explicitly via proof terms (specifically, cuts): in our system, a principal task of let-binding, a kind of cut, is to focus on the left (to process the list

of arguments in a bound function application); and a principal task of pattern matching, another kind of cut, is to focus on the right (to decompose the value being matched against a pattern).

From a Curry–Howard view, let-binding and pattern matching are different kinds of cuts. The cut formula A —basically, the type being matched or let-bound—must be *synthesized* (inferred) as an output (judgmentally, $\cdots \Rightarrow A$) from *heads* h (variables and annotated values) or *bound expressions* g (function application and annotated returner expressions); and ultimately, the outcomes of these cuts in our system are synthesized. But all other program terms are *checked* against input types A : judgmentally, $\cdots \Leftarrow A \cdots$ or $\cdots [A] \vdash \cdots$. In this sense, both our declarative and algorithmic type systems are *bidirectional* [Dunfield and Krishnaswami 2021].

In inversion stages, that is, expression typechecking (where a negative type is on the right of a turnstile) and pattern matching (where a positive type is on the left of a turnstile), refinements often need to be *extracted* from types in order to be used. For example, suppose we want to check the expression $\lambda x. \text{return } x$ against the type $(1 \wedge \text{ff}) \rightarrow \uparrow(1 \wedge \text{ff})$, which is semantically equivalent to $\text{ff} \supset 1 \rightarrow \uparrow(1 \wedge \text{ff})$. We need to extract ff (of $(1 \wedge \text{ff}) \rightarrow \cdots$) and put it in the logical context so that we can later use it (in typechecking the value x) to verify ff (of $\cdots \uparrow(1 \wedge \text{ff})$). If we instead put $x : 1 \wedge \text{ff}$ in the program context, then ff would not be usable (unless we can extract it from the program context, but it's simpler to extract from types as soon as possible rather than to extend the extraction judgment or to add another one) and typechecking would fail (it should succeed). Similarly, since subtyping may be viewed as implication, index information from positive subtypes or negative supertypes needs to be extracted for use. Declaratively, it is okay not to extract eagerly at polarity shifts in subtyping (the subtyping rules that extract are invertible), but it seems necessary in the algorithmic system.

Our declarative system includes two focusing stages, one (value typechecking) for positive types on the right of the turnstile (\vdash), and the other (spine typing) for negative types on the left. Our algorithmic system closely mirrors the declarative one, but does not conjure index instantiations or witnesses (like t in $\text{DeclSpine}\forall$ below), and instead systematically introduces and solves existential variables (like solving the existential variable \hat{a} as t in $\text{AlgSpine}\forall$ below), which we keep in algorithmic contexts Δ .

$$\frac{\Theta \vdash t : \tau \quad \Theta; \Gamma; [[t/a]N] \vdash s \gg \uparrow P}{\Theta; \Gamma; [\forall a : \tau. N] \vdash s \gg \uparrow P} \text{DeclSpine}\forall$$

$$\frac{\Theta; \Delta, \hat{a} : \tau; \Gamma; [[\hat{a}/a]N] \vdash s \gg \uparrow P / \chi \vdash \Delta', \hat{a} : \tau = t}{\Theta; \Delta; \Gamma; [\forall a : \tau. N] \vdash s \gg \uparrow P / \chi \vdash \Delta'} \text{AlgSpine}\forall$$

For example, applying a function of type $\forall b : \mathbb{N}. \text{List}(\text{Nat})(b) \rightarrow \cdots$ to the list $[4, 1, 2]$ should solve b to an index semantically equal to 3; the declarative system guesses an index term (like $3 + 0$), but the algorithmic system mechanically solves for it.

Our algorithmic right-focusing judgment has the form $\Theta; \Delta; \Gamma \vdash v \Leftarrow P / \chi \vdash \Delta'$, where χ is an output list of typing constraints and Δ' is an output context that includes index solutions to existentials. Similarly, the left-focusing stage is $\Theta; \Delta; \Gamma; [N] \vdash s \gg \uparrow P / \chi \vdash \Delta'$; it focuses on decomposing N (the type of the function being applied), introducing its existential variables for the arguments in the list s (sometimes called a *spine* [Cervesato and Pfenning 2003]), and outputting its guards to verify at the end of decomposition (an upshift). These existential variables flow to the right-focusing stage (value typechecking) and are solved there, possibly via subtyping. Constraints χ are only checked right at the end of focusing stages, when all their existential variables are solved.

For example, consider our rule for (fully) applying a function h to a list of arguments s :

$$\frac{\Theta; \Gamma \triangleright h \Rightarrow \downarrow N \quad \Theta; \cdot; \Gamma; [N] \vdash s \gg \uparrow P / \chi \dashv \cdot \quad \Theta; \Gamma \triangleleft \chi}{\Theta; \Gamma \triangleright h(s) \Rightarrow \uparrow P} \text{AlgSpineApp}$$

After synthesizing a thunk type $\downarrow N$ for the function h we are applying, we process the entire list of arguments s , until N 's return type $\uparrow P$. All (existential) value-determined indexes Ξ_N of N are guaranteed to be solved by the time an upshift is reached, and these solutions are eagerly applied to constraints χ , so that χ does not have existential variables and is hence SMT-solvable ($\Theta; \Gamma \triangleleft \chi$). The polarization of CBPV helps guarantee all solutions have no existential variables. Focusing stages introduce existential variables to input types, which may appear initially as a positive supertype in the subtyping premise for typechecking (value) variables. These existential variables are solved using the positive subtype, which never has existential variables. Dually, negative subtypes may have existential variables, but negative supertypes never do.

Our system requires intermediate computations like $h(s)$ to be explicitly named and sequenced by let-binding (a kind of *A-normal* [Flanagan et al. 1993] or *let-normal form*). Combined with focusing, this allows us to use (within the value typechecking stage) subtyping only in the typing rule for (value) variables. This makes our subsumption rule syntax-directed, simplifying and increasing the efficiency of our algorithm. We nonetheless prove a general subsumption lemma, which is needed to prove that substitution respects typing, a key syntactic or operational correctness property.

Focusing also gives us pattern matching for free [Krishnaswami 2009]: from a Curry–Howard view, the left-inversion stage is pattern matching. The (algorithmic³) left-inversion stage in our system is written $\Theta; \Gamma; [P] \triangleright \{r_i \Rightarrow e_i\}_{i \in I} \Leftarrow N$: it decomposes the positive P (representing the pattern being matched) to the left of the turnstile (written \triangleright to distinguish the algorithmic judgment from the corresponding declarative judgment, which instead uses \vdash). Our system is “half” strongly focused: we eagerly apply right-invertible but not left-invertible rules. This makes pattern matching in our system resemble the original presentation of pattern matching in CBPV. From a Curry–Howard view, increasing the strength of focusing would permit nested patterns.

A pattern type can have index equality constraints, such as for refined ADT constructors (for example, that the length of an empty list is zero) as output by unrolling. By using these equality constraints, we get a coverage-checking algorithm. For example, consider checking `get` (introduced in Sec. 1) against the type

$$\forall l, k : \mathbb{N}. \{v : \text{List } A \mid \text{len } v = l\} \rightarrow (\{v : \text{Nat} \mid \text{index } v = k\} \wedge (k < l)) \rightarrow \uparrow A$$

At the clause

$$\text{get } [] \ y = \text{unreachable}$$

we extract a logically inconsistent context ($l : \mathbb{N}, k : \mathbb{N}, l = 0, k < l$), which entails that `unreachable` checks against any type. Proof-theoretically, this use of equality resembles the elimination rule for Girard–Schroeder–Heister equality [Girard 1992; Schroeder-Heister 1994].

Bidirectional typing controls the flow of type information. Focusing in our system directs the flow of index information. The management of the flow of type refinement information, via the stratification of both focusing and bidirectionality, makes our algorithmic metatheory highly manageable.

3 EXAMPLE: VERIFYING MERGESORT

We show how to verify a non-structurally recursive mergesort function in our system: namely, that it terminates and returns a list with the same length as the input list (to verify it returns an ordered

³We give the algorithmic judgment to note existential variables Δ do not flow through it, or any of the non-focusing stages.

list, we need to extend our system with multi-argument measures, which is outside the scope of this paper). We only consider sorting lists of natural numbers Nat , defined as $\exists n : \mathbb{N}. \text{Nat}(n)$. For clarity, and continuity with Sections 1 and 2, we sometimes use syntactic sugar such as clausal pattern-matching, combining let-binding with pattern-matching on the let-bound variable, using “if-then-else” rather than pattern-matching on boolean type, and combining two or more pattern-matching expressions into one with a nested pattern such as $x :: y :: xs$.

Given type A and $n : \mathbb{N}$, we define $\text{List}(A)(n)$ by $\{v : \text{List } A \mid \text{len } v = n\}$. Modulo a small difference (see Sec. 4.6), our unrolling judgment unrolls $\text{List}(A)(n)$ to

$$(1 \wedge (n = 0)) + (A \times \exists n' : \mathbb{N}. \{v : \text{List } A \mid \text{len } v = n'\} \wedge (n = 1 + n'))$$

which is a refinement of $1 + (|A| \times \text{List } |A|)$. This is an unrolling of the inductive type, not the inductive type itself, so we must roll values of it into the inductive type. We use syntactic sugar: namely, $[]$ stands for $\text{into}(\text{inj}_1 ())$ and $x :: xs$ stands for $\text{into}(\text{inj}_2 \langle x, xs \rangle)$.

Just as we need a natural number type associating natural number program values with natural number indexes, we need a boolean type of values corresponding to boolean indexes. To this end, define the measure

$$\begin{aligned} \text{ixbool} &: (1 + 1) \rightarrow \mathbb{B} \\ \text{ixbool true} &= \text{tt} \\ \text{ixbool false} &= \text{ff} \end{aligned}$$

Given $b : \mathbb{B}$, the singleton type of boolean b is $\text{Bool}(b) = \{v : \text{Bool} \mid \text{ixbool } v = b\}$. Our unrolling judgment (Sec. 4.6) outputs the following type, a refinement of the boolean type encoded as $1 + 1$:

$$(1 \wedge (b = \text{tt})) + (1 \wedge (b = \text{ff}))$$

We encode true as $\text{into}(\text{inj}_1 ())$ which has type $\text{Bool}(\text{tt})$, and false as $\text{into}(\text{inj}_2 ())$ which has type $\text{Bool}(\text{ff})$. The boolean type Bool is defined as $\exists b : \mathbb{B}. \text{Bool}(b)$.

Assume we have the following:

$$\begin{aligned} \text{add} &: \downarrow (\forall m, n : \mathbb{N}. \text{Nat}(m) \rightarrow \text{Nat}(n) \rightarrow \uparrow \text{Nat}(m + n)) \\ \text{sub} &: \downarrow (\forall m, n : \mathbb{N}. (n \leq m) \supset \text{Nat}(m) \rightarrow \text{Nat}(n) \rightarrow \uparrow \text{Nat}(m - n)) \\ \text{div} &: \downarrow (\forall m, n : \mathbb{N}. (n \neq 0) \supset \text{Nat}(m) \rightarrow \text{Nat}(n) \rightarrow \uparrow \text{Nat}(m \div n)) \\ \text{lt} &: \downarrow (\forall m, n : \mathbb{N}. \text{Nat}(m) \rightarrow \text{Nat}(n) \rightarrow \uparrow \exists b : \mathbb{B}. \text{Bool}(b) \wedge (b = (m < n))) \\ \text{len} &: \downarrow (\forall n : \mathbb{N}. \text{List}(\text{Nat})(n) \rightarrow \uparrow \text{Nat}(n)) \\ [] &: \text{List}(\text{Nat})(0) \\ (::) &: \downarrow (\forall n : \mathbb{N}. \text{Nat} \rightarrow \text{List}(\text{Nat})(n) \rightarrow \uparrow \text{List}(\text{Nat})(1 + n)) \end{aligned}$$

The SMT solver Z3 [de Moura and Bjørner 2008], for example, supports integer division (and modulo and remainder operators); internally, these operations are translated to multiplication. Here, we are considering natural number indexes, but we can add the constraint $n \geq 0$ (for naturals n) when translating them to integers in an SMT solver such as Z3. Allowing integer division in general is not SMT decidable, but for this example, n is always instantiated to a constant (2), which is decidable. Note that Z3 supports division by zero, but our div has a guard requiring the divisor to be nonzero ($n \neq 0$), so we need not consider this. Division on naturals takes the floor of what would otherwise be a rational number (for example, $3 \div 2 = 1$).

First, we define the function merge for merging two lists while preserving order. It takes two lists as inputs, but also a natural number representing the sum of their lengths. Since at least one list decreases in length at recursive calls, so does the sum of their lengths, implying the function terminates when applied. However, in the *refined* system presented in (the figures of) this paper, to keep things simple, we provide only one rule for recursive expressions, whose termination metric is $<$ on natural numbers. Because the system as presented lacks a rule that supports a termination

metric on sums $n_1 + n_2$ of natural numbers, we need to reify the sum $n = n_1 + n_2$ of the length indexes n_1 and n_2 into a ghost parameter n of (asserting) singleton type. However, we emphasize the ghost parameter in this example is not a fundamental limitation of our system, because our system can be extended to include other termination metrics such as $<$ on the sum of natural numbers (we discuss this further in Sec. 4.7).

```

merge : ∀n, n1, n2 : ℕ. Nat(n) ∧ (n = n1 + n2) → List(Nat)(n1) → List(Nat)(n2)
      → ↑List(Nat)(n1 + n2)
merge y [] xs2 = return xs2
merge y xs1 [] = return xs1
merge (succ y) (x1 :: xs1) (x2 :: xs2) =
  if lt(x1, x2) then
    let recresult = merge(y, xs1, (x2 :: xs2));
    let result = x1 :: recresult;
    return result
  else
    let recresult = merge(y, (x1 :: xs1), xs2);
    let result = x2 :: recresult;
    return result

```

In a well-typed let-binding $\text{let } x = g; e$ the bound expression g is a value-returning computation (that is, has upshift type), and e is a computation that binds to x the value (of positive type) resulting from computing g . (Liquid Haskell, lacking CBPV’s type distinction between computations and values, instead approximates whether binders terminate to a value.) Since x has positive type, we can match it against patterns (see, for example, the final clause of `split`, discussed next).

We now define the function `split` that takes a list and splits it into two lists. It is a standard “every-other” implementation, and we have to be a bit careful about the type refinement so as not to be “off by one” in the lengths of the resulting lists.

```

split : ∀n : ℕ. List(Nat)(n) → ↑(List(Nat)((n + 1) ÷ 2) × (List(Nat)(n - ((n + 1) ÷ 2)))
split [] = return ⟨[], []⟩
split [x] = return ⟨[x], []⟩
split x1 :: x2 :: xs =
  let recresult = split(xs);
  match recresult {
    ⟨xs1, xs2⟩ ⇒ return ⟨x1 :: xs1, x2 :: xs2⟩
  }

```

We are ready to implement `mergesort`, but since we use a ghost parameter, we need to define an auxiliary function `auxmergesort` additionally taking the length of the list being ordered. We

introduce syntactic sugar for a let-binding followed by pattern-matching on its result.

```

auxmergesort : ∀n : ℕ. Nat(n) → List(Nat)(n) → ↑List(Nat)(n)
auxmergesort y [] = return []
auxmergesort y [x] = return [x]
auxmergesort y xs =
  let ⟨leftxs, rightxs⟩ = split(xs);
  let lenleftxs = div(succ y, 2);
  let lenrightxs = y - div(succ y, 2);
  let sortleftxs = auxmergesort(lenleftxs, leftxs);
  let sortrightxs = auxmergesort(lenrightxs, rightxs);
  return merge(y, sortleftxs, sortrightxs)

```

Finally, we define a mergesort that is verified to terminate and to return a list of the same length as the input list.

```

mergesort : ∀n : ℕ. List(Nat)(n) → ↑List(Nat)(n)
mergesort xs =
  let len = len(xs);
  return mergesortaux(len, xs)

```

In the system of this paper, we cannot verify that mergesort returns a *sorted* list. This is because our system lacks multi-argument measures which can specify relations between indexes of different parts of a data type. (To handle this, we are extending our system with multi-argument measures, which is nontrivial and requires a significant degree of additional machinery outside the scope of this paper.) But this example is interesting nonetheless, because auxmergesort is *not* structurally recursive: its recursive calls are on lists obtained by splitting the input list roughly in half, not on the structure of the list ($- :: -$). Further, it showcases the main features of our foundational system, the declarative specification of which we turn to next, in Sec. 4.

4 DECLARATIVE SYSTEM

We present our core declarative calculus and type system.

First, we discuss the syntax of program terms, types, index terms, sorts, functors, algebras, and contexts. Then, in Sec. 4.1, we discuss the index sorting judgment $\Theta \vdash t : \tau$ and the propositional validity judgment $\Theta \vdash \phi$ true, index-level substitution, and basic logical properties required of the index domain. In Sec. 4.2, we discuss the well-formedness of (logical and program) contexts (Θ ctx and $\Theta \vdash \Gamma$ ctx), types ($\Theta \vdash A$ type $[\Xi]$), functors ($\Theta \vdash \mathcal{F}$ functor $[\Xi]$), and algebras ($\Xi; \Theta \vdash \alpha : F(\tau) \Rightarrow \tau$). In Sec. 4.3, we discuss judgmental equivalence. In Sec. 4.4, we discuss extraction ($\Theta \vdash A \rightsquigarrow A' [\Theta']$). In Sec. 4.5, we discuss subtyping ($\Theta \vdash A \leq B$). In Sec. 4.6, we discuss the unrolling judgment for refined inductive types. In Sec. 4.7, we discuss the typing system. In Sec. 4.8, we extend substitution to that of program values for program variables, and prove a substitution lemma stating that typing is stable under substitution (a key operational correctness result).

In Fig. 1, we summarize the key judgments constituting the declarative system. In the figure, “pre.” abbreviates “presupposes”, which basically lists the judgments (rightmost column) we tend to leave implicit in rules defining the given judgment (leftmost column). Presuppositions also indicate the (input or output) moding of judgments. For example, on the one hand $\Theta; \Gamma \vdash v \Leftarrow P$ presupposes $\Theta \vdash \Gamma$ ctx and $\Theta \vdash P$ type $[\Xi_P]$ for some Ξ_P , where the former presupposes Θ ctx, and Θ, Γ , and P

are input-moded; on the other hand, $\Theta; \Gamma \vdash h \Rightarrow P$ does not presuppose $\Theta \vdash P \text{ type}[\Xi_P]$ for some Ξ_P , but rather we must prove that the output-moded P is well-formed (which is straightforward). Groups of mutually defined judgments are separated by blank lines.

$\Theta \text{ ctx}$	(Sec. 4.2)	pre.	no judgment
$\Theta \vdash t : \tau$	(Sec. 4.1)	pre.	$\Theta \text{ ctx}$
$\Theta \vdash \phi \text{ true}$	(Sec. 4.1)	pre.	$\Theta \vdash \phi : \mathbb{B}$
$\Theta \vdash A \text{ type}[\Xi]$	(Sec. 4.2)	pre.	$\Theta \text{ ctx}$
$\Theta \vdash \mathcal{F} \text{ functor}[\Xi]$	(Sec. 4.2)	pre.	$\Theta \text{ ctx}$
$\Xi; \Theta \vdash \alpha : F(\tau) \Rightarrow \tau$	(Sec. 4.2)	pre.	$\Xi \text{ ctx}$ and $\Theta \vdash F \text{ functor}[\Xi_F]$
$\Theta \vdash \mathcal{F} \equiv \mathcal{G}$	(Sec. 4.3)	pre.	$\Theta \vdash \mathcal{F} \text{ functor}[\Xi_{\mathcal{F}}]$ and $\Theta \vdash \mathcal{G} \text{ functor}[\Xi_{\mathcal{G}}]$
$\Theta \vdash A \equiv^{\pm} B$	(Sec. 4.3)	pre.	$\Theta \vdash A \text{ type}[\Xi_A]$ and $\Theta \vdash B \text{ type}[\Xi_B]$
$\Theta \vdash A \rightsquigarrow^{\pm} A' [\Theta_A]$	(Sec. 4.4)	pre.	$\Theta \vdash A \text{ type}[\Xi_A]$
$\Theta \vdash A \leq^{\pm} B$	(Sec. 4.5)	pre.	$\Theta \vdash A \text{ type}[\Xi_A]$ and $\Theta \vdash B \text{ type}[\Xi_B]$
$\Xi; \Theta \vdash \text{unroll}_{F;\alpha}(G; \beta; \tau; t) \doteq P$	(Sec. 4.6)	pre.	$\Xi; \Theta \vdash \alpha : F(\tau) \Rightarrow \tau$ and $\Xi; \Theta \vdash \beta : G(\tau) \Rightarrow \tau$ and $\Theta \vdash t : \tau$
$\Theta \vdash \Gamma \text{ ctx}$	(Sec. 4.2)	pre.	$\Theta \text{ ctx}$
$\Theta; \Gamma \vdash h \Rightarrow P$	(Sec. 4.7)	pre.	$\Theta \vdash \Gamma \text{ ctx}$
$\Theta; \Gamma \vdash g \Rightarrow \uparrow P$	(Sec. 4.7)	pre.	$\Theta \vdash \Gamma \text{ ctx}$
$\Theta; \Gamma \vdash v \Leftarrow P$	(Sec. 4.7)	pre.	$\Theta \vdash \Gamma \text{ ctx}$ and $\Theta \vdash P \text{ type}[\Xi_P]$
$\Theta; \Gamma \vdash e \Leftarrow N$	(Sec. 4.7)	pre.	$\Theta \vdash \Gamma \text{ ctx}$ and $\Theta \vdash N \text{ type}[\Xi_N]$
$\Theta; \Gamma; [P] \vdash \{r_i \Rightarrow e_i\}_{i \in I} \Leftarrow N$	(Sec. 4.7)	pre.	$\Theta \vdash \Gamma \text{ ctx}$ and $\Theta \vdash P \text{ type}[\Xi_P]$ and $\Theta \vdash N \text{ type}[\Xi_N]$
$\Theta; \Gamma; [N] \vdash s \gg \uparrow P$	(Sec. 4.7)	pre.	$\Theta \vdash \Gamma \text{ ctx}$ and $\Theta \vdash N \text{ type}[\Xi_N]$

Fig. 1. Declarative judgments and their presuppositions

Program terms. Program terms are defined in Fig. 2. We polarize terms into two main syntactic categories: expressions (which have negative type) and values (which have positive type). Program terms are further distinguished according to whether their principal types are synthesized (heads and bound expressions) or checked (spines and patterns).

Expressions e consist of functions $\lambda x. e$, recursive expressions $\text{rec } x : N. e$, let-bindings $\text{let } x = g; e$, match expressions $\text{match } h \{r_i \Rightarrow e_i\}_{i \in I}$, value returners (or producers) $\text{return } v$, and an unreachable expression unreachable (such as an impossible match). Bound expressions g , which can be let-bound, consist of expressions annotated with a returner type ($e : \uparrow P$) and applications $h(s)$ of a head h to a spine s . Heads h , which can be applied to a spine or pattern-matched, consist of variables x and positive-type-annotated values ($v : P$). Spines s are lists of values; we often omit the empty spine \cdot , writing (for example) v_1, v_2 instead of v_1, v_2, \cdot . In match expressions, heads are matched against patterns r .

Values consist of variables x , the unit value $\langle \rangle$, pairs $\langle v_1, v_2 \rangle$, injections into sum type $\text{inj}_k v$ where k is 1 or 2, rollings into inductive type $\text{into}(v)$, and thunks (suspended computations) $\{e\}$.

Program variables	x, y, z
Expressions	$e ::= \text{return } v \mid \text{let } x = g; e \mid \text{match } h \{r_i \Rightarrow e_i\}_{i \in I} \mid \lambda x. e \mid \text{rec } x : N. e \mid \text{unreachable}$
Values	$v ::= x \mid \langle \rangle \mid \langle v, v \rangle \mid \text{inj}_1 v \mid \text{inj}_2 v \mid \text{into}(v) \mid \{e\}$
Heads	$h ::= x \mid (v : P)$
Bound expressions	$g ::= h(s) \mid (e : \uparrow P)$
Spines	$s ::= \cdot \mid v, s$
Patterns	$r ::= \text{into}(x) \mid \langle \rangle \mid \langle x, y \rangle \mid \text{inj}_1 x \mid \text{inj}_2 x$

Fig. 2. Program terms

Types. Types are defined in Fig. 3. Types are polarized into positive (value) types P and negative (computation) types N . We write A, B and C for types of either polarity.

Positive types	$P, Q ::= 1 \mid P \times Q \mid 0 \mid P + Q \mid \downarrow N \mid \{v : \mu F \mid (\text{fold}_F \alpha) v =_\tau t\} \mid \exists a : \tau. P \mid P \wedge \phi$
Negative types	$N, M ::= P \rightarrow N \mid \uparrow P \mid \forall a : \tau. N \mid \phi \supset N$
Types	$A, B, C ::= P \mid N$

Fig. 3. Types

Positive types consist of the unit type 1 , products $P_1 \times P_2$, the void type 0 , sums $P_1 + P_2$, downshifts (of negative types; *thunk* types) $\downarrow N$, *asserting* types $P \wedge \phi$ (read “ P with ϕ ”), index-level existential quantifications $\exists a : \tau. P$, and refined inductive types $\{v : \mu F \mid (\text{fold}_F \alpha) v =_\tau t\}$. We read $\{v : \mu F \mid (\text{fold}_F \alpha) v =_\tau t\}$ as the type having values v of inductive type μF (with signature F) such that the (index-level) *measurement* $(\text{fold}_F \alpha) v =_\tau t$ holds; in Sec. 4.0.1 and Sec. 5, we explain the metavariables F , α , τ , and t , as well as what these and the syntactic parts μ and fold denote. Briefly, μ roughly denotes “least fixed point of” and a fold over F with α (having carrier sort τ) indicates a measure on the inductive type μF into τ .

Negative types consist of function types $P \rightarrow N$, upshifts (of positive types; *lift* or *returning* types) $\uparrow P$ (dual to $\downarrow N$), propositionally *guarded* types $\phi \supset N$ (read “ ϕ implies N ”; dual to $P \wedge \phi$), and index-level universal quantifications $\forall a : \tau. N$ (dual to $\exists a : \tau. P$).

In $P \wedge \phi$ and $\phi \supset N$, the index proposition ϕ has no run-time content. Neither does the a in $\exists a : \tau. P$ and $\forall a : \tau. N$, nor the recursive refinement predicate $(\text{fold}_F \alpha) v =_\tau t$ in $\{v : \mu F \mid (\text{fold}_F \alpha) v =_\tau t\}$.

Index language: sorts, terms, and propositions. Our type system is parametric in the index domain, provided the latter has certain (basic) properties. For our system to be decidable, the index domain must be decidable. It is instructive to work with a specific index domain: Figure 4 defines a quantifier-free logic of linear equality, inequality, and arithmetic, which is decidable [Barrett et al. 2009].

Sorts τ consist of booleans \mathbb{B} , natural numbers \mathbb{N} , integers \mathbb{Z} , and products $\tau_1 \times \tau_2$. Index terms t consist of variables a , numeric constants n , addition $t_1 + t_2$, subtraction $t_1 - t_2$, pairs (t_1, t_2) , projections $\pi_1 t$ and $\pi_2 t$, and propositions ϕ . Propositions ϕ (the logic of the index domain) are built over index terms, and consist of equality $t_1 = t_2$, inequality $t_1 \leq t_2$, conjunction $\phi_1 \wedge \phi_2$, disjunction $\phi_1 \vee \phi_2$, negation $\neg \phi$, trivial truth tt , and trivial falsity ff .

Sorts	$\tau ::= \mathbb{B} \mid \mathbb{N} \mid \mathbb{Z} \mid \tau \times \tau$
Index variables	a, b, c
Index terms	$t ::= a \mid n \mid t + t \mid t - t \mid (t, t) \mid \pi_1 t \mid \pi_2 t \mid \phi$
Propositions	$\phi, \psi ::= t = t \mid t \leq t \mid \phi \wedge \phi \mid \phi \vee \phi \mid \neg \phi \mid \text{tt} \mid \text{ff}$

Fig. 4. Index domain

4.0.1 Inductive types, functors, and algebras. We encode algebraic data types (and measures on them) using their standard semantics. In the introduction (Sec. 1), to refine the type of A -lists by their length, we defined a recursive function `len` over the inductive structure of lists. Semantically, we characterize this structural recursion by algebraic folds over polynomial endofunctors; we design our system in line with this semantics. While this presentation may appear overly abstract for the user, it should be possible to allow the user to use the same or similar syntax as programs to express measures if they annotate them as measures in the style of Liquid Haskell.

We express inductive type structure without reference to constructor names by syntactic functors resembling the polynomial functors. For example (modulo the difference for simplifying unrolling), we can specify the signature of the inductive type of lists of terms of type A syntactically as a functor $\underline{1} \oplus (\underline{A} \otimes \text{Id})$, where \underline{C} denotes the constant (set) functor (sending any set to the set denoted by type C), Id denotes the identity functor (sending any set to itself), the denotation of $F_1 \otimes F_2$ sends a set X to the product $(\llbracket F_1 \rrbracket X) \times (\llbracket F_2 \rrbracket X)$, and the denotation of $F_1 \oplus F_2$ sends a set X to the disjoint union $(\llbracket F_1 \rrbracket X) \uplus (\llbracket F_2 \rrbracket X)$. The idea is that each component of the sum functor \oplus represents a data constructor, so that (for example) $\underline{1}$ represents the nullary constructor $[]$, and \underline{A} represents the head element of a cons cell which is attached (via \otimes) to the recursive tail list represented by Id .

A functor F (Fig. 5) is a sum (\oplus) of products (\hat{P}), which multiply (\otimes) base functors (B) consisting of identity functors that represent recursive positions (Id) and constant functors (\underline{P}) at positive type P . The rightmost factor in a product \hat{P} is the (product) unit functor I . By convention, \otimes has higher precedence than \oplus . For convenience in specifying functor well-formedness (appendix Fig. 7) and denotation (appendix Fig. 37), \mathcal{F} is a functor F or a base functor B .

Functors	$F, G, H ::= \hat{P} \mid F \oplus F$
	$\hat{P} ::= I \mid B \otimes \hat{P}$
	$B ::= \underline{P} \mid \text{Id}$
	$\mathcal{F} ::= F \mid B$

Fig. 5. Functors

A direct grammar F for sums of products (of constant and identity functors) consists of $F ::= \hat{P} \mid F \oplus F$ and $\hat{P} ::= B \mid B \otimes \hat{P}$ and $B ::= \underline{P} \mid \text{Id}$. The grammar $F ::= F \oplus F \mid F \otimes F \mid B$ is semantically equivalent to sums of products, but syntactically inconvenient, because it allows writing products of sums. We do not use either of these grammars, but rather Fig. 5, because it simplifies inductive type unrolling (Sec. 4.6). (In any case, a surface language where data types have named constructors would have to be elaborated to use one of these grammars.) These grammars have naturally isomorphic interpretations. For example, in our functor grammar (Fig. 5), we instead write $\text{NatF} = I \oplus (\text{Id} \otimes I)$ (note I is semantically equivalent to $\underline{1}$): notice that for any set X , we have $\llbracket I \oplus (\text{Id} \otimes I) \rrbracket X = 1 \uplus (X \times 1) \cong 1 \uplus X = \llbracket \underline{1} \oplus \text{Id} \rrbracket X$.

As we will discuss in Sec. 5, every polynomial endofunctor F has a fixed point μF satisfying a recursion principle for defining measures (on μF) by folds with algebras. We define algebras in Fig. 6. An algebra α is a list of clauses $p \Rightarrow t$ which pattern match on algebraic structure (p , q , and o are patterns) and bind variables in index bodies t . Sum algebra patterns p consist of $\text{inj}_1 p$ and $\text{inj}_2 p$ (which match on sum functors \oplus). Product algebra patterns q consist of tuples (o, q) (which match on \otimes) ending in the unit pattern $()$ (which match on I). Base algebra patterns o consist of wildcard patterns \top (which match on constant functors \underline{P}), variable patterns a (which match on the identity functor Id), and pack patterns $\text{pack}(a, o)$ (which match on *existential* constant functors $\exists a : \tau. P$, where a is also bound in the bodies t of algebra clauses).

Algebras	$\alpha, \beta ::= \cdot \mid (p \Rightarrow t \mid \alpha)$
Sum algebra patterns	$p ::= \text{inj}_1 p \mid \text{inj}_2 p \mid q$
Product algebra patterns	$q ::= () \mid (o, q)$
Base algebra patterns	$o ::= \top \mid a \mid \text{pack}(a, o)$

Fig. 6. Algebras

For example, given a type P , consider the functor $I \oplus (\underline{P} \otimes \text{Id} \otimes I)$. To specify the function $\text{length} : \text{List } P \rightarrow \mathbb{N}$ computing the length of a list of values of type P , we write the algebra $\text{inj}_1 () \Rightarrow 0 \mid \text{inj}_2 (\top, (a, ())) \Rightarrow 1 + a$ with which to fold $\text{List } P$.

With the pack algebra pattern, we can use indexes of an inductive type in our measures. For example, given $a : \mathbb{N}$, and defining the singleton type $\text{Nat}(a)$ as $\{v : \mu \text{NatF} \mid (\text{fold}_{\text{NatF}} \text{ixnat}) v = a\}$ where $\text{NatF} = I \oplus \text{Id} \otimes I$ and $\text{ixnat} = \text{inj}_1 () \Rightarrow 0 \mid \text{inj}_2 (a, ()) \Rightarrow 1 + a$, consider lists of natural numbers, specified by $I \oplus \exists b : \mathbb{N}. \text{Nat}(b) \otimes \text{Id} \otimes I$. Folding such a list with the algebra $\text{inj}_1 () \Rightarrow 0 \mid \text{inj}_2 (\text{pack}(b, \top), a, ()) \Rightarrow a + b$ sums all the numbers in the list. (For clarity, we updated the definitions in Sec. 2 to agree with our grammars as presented in Fig. 5 and Fig. 6.)

For measures relating indexes in structurally distinct positions within an inductive type, in ongoing work we are extending our system with multi-argument measures by way of higher-order sorts $\tau_1 \Rightarrow \tau_2$. Doing so would allow us to refine, for example, integer lists to lists of integers *in increasing order*, because we could then compare the indexed elements of a list.

Contexts. A *logical context* $\Theta ::= \cdot \mid \Theta, a : \tau \mid \Theta, \phi$ is an ordered list of index propositions ϕ and (index) variable sortings $a : \tau$ (which may be used in subsequent propositions). A *program variable context* (or *program context*) $\Gamma ::= \cdot \mid \Gamma, x : P$ is a set of (program) variable typings $x : P$. A value-determined context $\Xi ::= \cdot \mid \Theta, a : \tau$ is a set of index sortings $a : \tau$. In any kind of context, a variable can be declared at most once.

4.1 Index sorting and propositional validity

We have a standard index sorting judgment $\Theta \vdash t : \tau$ (appendix Fig. 4) checking that, under context Θ , index term t has sort τ . For example, $a : \mathbb{N} \vdash \neg(a \leq a + 1) : \mathbb{B}$. This judgment does not depend on propositions in Θ , which only matter when checking propositional validity (only done in subtyping and program typing). The operation $\overline{\Theta}$ merely removes all propositions ϕ from Θ .

Well-sorted index terms $\Theta \vdash t : \tau$ denote functions $\llbracket t \rrbracket : \llbracket \overline{\Theta} \rrbracket \rightarrow \llbracket \tau \rrbracket$. For each Θ , define $\llbracket \Theta \rrbracket$ as the set of index-level semantic substitutions (defined in this paragraph) $\{\delta \mid \vdash \delta : \Theta\}$. For example, $\llbracket 4 + a \rrbracket_{3/a} = 7$ and $\llbracket b = 1 + 0 \rrbracket_{1/b} = \{\bullet\}$ (that is, true) and $\llbracket a = 1 \rrbracket_{2/a} = \emptyset$ (that is, false). An *index-level semantic substitution* $\vdash \delta : \Theta$ assigns exactly one semantic index value d to each index variable

in $\text{dom}(\Theta)$ such that every proposition ϕ in Θ is true (written $\{\bullet\}$; false is \emptyset):

$$\frac{}{\vdash \cdot : \cdot} \quad \frac{\vdash \delta : \Theta \quad d \in \llbracket \tau \rrbracket \quad a \notin \text{dom}(\Theta)}{\vdash (\delta, d/a) : (\Theta, a : \tau)} \quad \frac{\vdash \delta : \Theta \quad \llbracket \phi \rrbracket_{\delta} = \{\bullet\}}{\vdash \delta : (\Theta, \phi)}$$

A *propositional validity* or *truth* judgment $\Theta \vdash \phi$ true, which is a *semantic* entailment relation, holds if ϕ is valid under Θ , that is, if ϕ is true under every interpretation of variables in Θ such that all propositions in Θ are true. We say t and t' are *logically equal* under Θ if $\Theta \vdash t = t'$ true.

An *index-level syntactic substitution* σ is a list of index terms to be substituted for index variables: $\sigma ::= \cdot \mid \sigma, t/a$. The metaoperation $[\sigma]O$, where O is an index term, program term, or type, is sequential substitution: $[\cdot]O = O$ and $[\sigma, t/a]O = [\sigma]([t/a]O)$, where $[t/a]O$ is standard capture-avoiding (by α -renaming) substitution. Syntactic substitutions (index-level) are typed (“sorted”) in a standard way. Because syntactic substitutions substitute terms that may have free variables, their judgment form includes a context to the left of the turnstile, in contrast to semantic substitution:

$$\frac{}{\Theta_0 \vdash \cdot : \cdot} \quad \frac{\Theta_0 \vdash \sigma : \Theta \quad \Theta_0 \vdash [\sigma]t : \tau \quad a \notin \text{dom}(\Theta)}{\Theta_0 \vdash (\sigma, t/a) : (\Theta, a : \tau)} \\ \frac{\Theta_0 \vdash \sigma : \Theta \quad \Theta_0 \vdash [\sigma]\phi \text{ true}}{\Theta_0 \vdash \sigma : (\Theta, \phi)}$$

Because our substitution operation $[\sigma]$ – applies sequentially, we type (“sort”) the application of the rest of the substitution to the head being substituted. For example, the rule concluding $\Theta_0 \vdash (\sigma, t/a) : (\Theta, a : \tau)$ checks that the application $[\sigma]t$ of σ to t has sort τ .

The decidability of our system depends on the decidability of propositional validity. Our example index domain is decidable [Barrett et al. 2009]. Our system is parametric in the index domain, provided the latter has certain properties. In particular, propositional validity must satisfy the following basic properties required of a logical theory (Θ ctx is logical context well-formedness).

- *Weaken*: If $\Theta_1, \Theta, \Theta_2$ ctx and $\Theta_1, \Theta_2 \vdash \phi$ true, then $\Theta_1, \Theta, \Theta_2 \vdash \phi$ true.
- *Permute*: If Θ, Θ_1 ctx and Θ, Θ_2 ctx and $\Theta, \Theta_1, \Theta_2, \Theta' \vdash \phi$ true, then $\Theta, \Theta_2, \Theta_1, \Theta' \vdash \phi$ true.
- *Substitution*: If $\Theta \vdash \phi$ true and $\Theta_0 \vdash \sigma : \Theta$, then $\Theta_0 \vdash [\sigma]\phi$ true.
- *Equivalence*: The relation $\Theta \vdash t_1 = t_2$ true is an equivalence relation.
- *Assumption*: If Θ_1, ϕ, Θ_2 ctx, then $\Theta_1, \phi, \Theta_2 \vdash \phi$ true.
- *Consequence*: If $\Theta_1 \vdash \psi$ true and $\Theta_1, \psi, \Theta_2 \vdash \phi$ true, then $\Theta_1, \Theta_2 \vdash \phi$ true.
- *Consistency*: It is not the case that $\cdot \vdash \text{ff}$ true.

We also assume that $\Theta \vdash t : \tau$ is decidable and satisfies weakening and substitution. Our example index domain satisfies all these properties.

4.2 Well-formedness

Context well-formedness Θ ctx and $\Theta \vdash \Gamma$ ctx (appendix Fig. 8) is straightforward. For both logical and program context well-formedness, there can be at most one of each variable. Index terms in well-formed logical contexts must have boolean sort:

$$\frac{}{\cdot \text{ ctx} \text{ LogCtxEmpty}} \quad \frac{\Theta \text{ ctx} \quad a \notin \text{dom}(\Theta)}{(\Theta, a : \tau) \text{ ctx} \text{ LogCtxVar}} \quad \frac{\Theta \text{ ctx} \quad \Theta \vdash \phi : \mathbb{B}}{(\Theta, \phi) \text{ ctx} \text{ LogCtxProp}}$$

In well-formed program variable contexts $\Theta \vdash \Gamma$ ctx, the types (of program variables) must be well-formed under Θ ; further, we must not be able to extract index information from these types (in the sense of Sec. 4.4). For example, $x : 1 \wedge \text{ff}$ is an ill-formed program context because ff can be extracted, but $x : \downarrow \uparrow 1 \wedge \text{ff}$ is well-formed because nothing under a shift type can be extracted.

Type well-formedness $\Theta \vdash A \text{ type}[\Xi]$ (read “under Θ , type A is well-formed with value-determined indexes Ξ ”) has Ξ in output mode, which tracks index variables appearing in the type A that are uniquely⁴ determined by values of refined inductive types in A , particularly by their folds. (See Lemma 5.4 in Sec. 5.) Consider the following type well-formedness rule:

$$\frac{\Theta \vdash F \text{ functor}[\Xi] \quad ; \Theta \vdash \alpha : F(\tau) \Rightarrow \tau \quad (b : \tau) \in \Theta}{\Theta \vdash \{v : \mu F \mid (\text{fold}_F \alpha) v =_\tau b\} \text{ type}[\Xi \cup b : \tau]} \text{DeclTp}\mu\text{Var}$$

The index b is uniquely determined by a value of the conclusion type, so we add it to Ξ . For example, the value $\text{one} = \text{into}(\text{inj}_2 \langle \text{into}(\text{inj}_1 \langle \rangle), \langle \rangle \rangle)$ determines the variable b appearing in the value’s type $\text{NatF}(b) = \{v : \mu \text{NatF} \mid (\text{fold}_{\text{NatF}} \text{ixnat}) v = b\}$ to be one. (We have a similar rule where the b -position metavariable is not an index variable, adding nothing to Ξ .) We use set union ($\Xi \cup b : \tau$) as b may already be value-determined in F (that is, $(b : \tau)$ may be in Ξ). The algebra well-formedness premise $;\Theta \vdash \alpha : F(\tau) \Rightarrow \tau$ requires the algebra α to be closed (that is, the first context is empty, \cdot). This premise ensures that existential variables never appear in algebras, which is desirable because folds with algebras solve existential variables when typechecking a value (see Sec. 6).

Because ultimately Ξ tracks only measure-determined indexes, $\text{DeclTp}\mu\text{Var}$ is the only rule that adds to Ξ . The index propositions of asserting and guarded types do not track anything beyond what is tracked by the types to which they are connected.

$$\frac{\Theta \vdash P \text{ type}[\Xi] \quad \Theta \vdash \phi : \mathbb{B}}{\Theta \vdash P \wedge \phi \text{ type}[\Xi]} \text{DeclTp}\wedge \qquad \frac{\Theta \vdash N \text{ type}[\Xi] \quad \Theta \vdash \phi : \mathbb{B}}{\Theta \vdash \phi \supset N \text{ type}[\Xi]} \text{DeclTp}\supset$$

We restrict quantification to value-determined index variables in order to guarantee we can always solve them algorithmically. For example, in checking one against the type $\exists a : \mathbb{N}. \text{Nat}(a)$, we solve a to an index semantically equal to $1 \in \mathbb{N}$. If $\Theta, a : \tau \vdash P \text{ type}[\Xi]$, then $\exists a : \tau. P$ is well-formed if and only if $(a : \tau) \in \Xi$, and similarly for universal quantification (which we’ll restrict to the argument types of function types; argument types are positive):

$$\frac{\Theta, a : \tau \vdash P \text{ type}[\Xi, a : \tau]}{\Theta \vdash \exists a : \tau. P \text{ type}[\Xi]} \text{DeclTp}\exists \qquad \frac{\Theta, a : \tau \vdash N \text{ type}[\Xi, a : \tau]}{\Theta \vdash \forall a : \tau. N \text{ type}[\Xi]} \text{DeclTp}\forall$$

We read commas in value-determined contexts such as Ξ_1, Ξ_2 as set union $\Xi_1 \cup \Xi_2$ together with the fact that $\text{dom}(\Xi_1) \cap \text{dom}(\Xi_2) = \emptyset$, so these rules can be read top-down as removing a .

A value of product type is a pair of values, so we take the union of what each component value determines:

$$\frac{\Theta \vdash P_1 \text{ type}[\Xi_1] \quad \Theta \vdash P_2 \text{ type}[\Xi_2]}{\Theta \vdash P_1 \times P_2 \text{ type}[\Xi_1 \cup \Xi_2]} \text{DeclTp}\times$$

We also take the union for function types $P \rightarrow N$, because to use a function, due to focusing, we must provide values for all its arguments. The Ξ of $\text{Nat}(a) \rightarrow \uparrow \text{Nat}(a)$ is $a : \mathbb{N}$, so $\forall a : \mathbb{N}. \text{Nat}(a) \rightarrow \uparrow \text{Nat}(a)$ is well-formed. In applying a head of this type to a value, we must instantiate a to an index semantically equal to what that value determines; for example, if the value is one, then a gets instantiated to an index semantically equal to $1 \in \mathbb{N}$.

However, a value of sum type is either a left- or right-injected value, but we don’t know which, so we take the intersection of what each injection determines:

$$\frac{\Theta \vdash P_1 \text{ type}[\Xi_1] \quad \Theta \vdash P_2 \text{ type}[\Xi_2]}{\Theta \vdash P_1 + P_2 \text{ type}[\Xi_1 \cap \Xi_2]} \text{DeclTp}+$$

⁴Semantically speaking.

The unit type 1 and void (empty) type 0 both have empty Ξ . We also empty out value-determined indexes at shifts, preventing certain quantifications over shifts. For example, $\forall a : \mathbb{N}. \uparrow \text{Nat}(a)$ (which is void anyway) is not well-formed. Crucially, we will see that this restriction, together with focusing, guarantees indexes will be algorithmically solved by the end of certain stages.

We define functor and algebra well-formedness in Fig. 7 of the appendix.

Functor well-formedness $\Theta \vdash \mathcal{F}$ functor $[\Xi]$ is similar to type well-formedness: constant functors output the Ξ of the underlying positive type, the identity and unit functors Id and I have empty Ξ , the product functor $B \otimes \hat{P}$ takes the union of the component Ξ s, and the sum functor $F_1 \oplus F_2$ takes the intersection. The latter two reflect the fact that unrolling inductive types (Sec. 4.6) generates $+$ types from \oplus functors and \times types from \otimes functors. That I has empty Ξ reflects that 1 (unrolled from I) does too, together with the fact that asserting and guarded types do not affect Ξ .

Algebra well-formedness $\Xi; \Theta \vdash \alpha : F(\tau) \Rightarrow \tau$ (read “under Ξ and Θ , algebra α is well-formed and has type $F(\tau) \Rightarrow \tau$ ”) has two contexts: Ξ is for α (in particular, the index bodies of its clauses) and Θ is for F (in particular, the positive types of constant functors); we maintain the invariant that $\Xi \subseteq \Theta$. We have these separate contexts to prevent existential variables from appearing in α (as explained with respect to $\text{DeclTp}\mu\text{Var}$) while still allowing them to appear in F . For example, consider $\exists b : \mathbb{N}. \{v : \mu F(b) \mid (\text{fold}_{F(b)} \alpha) v = n\}$ where $F(b) = \underline{\text{Nat}(b)} \otimes I \oplus \underline{\text{Nat}(b)} \otimes \text{Id} \otimes I$ and $\alpha = \text{inj}_1(\top, ()) \Rightarrow 0 \mid \text{inj}_2(\top, a, ()) \Rightarrow 1 + a$.

Refined inductive type well-formedness initializes the input Ξ to \cdot , but index variables can be bound in the body of an algebra:

$$\frac{\Xi, a : \tau; \Theta, a : \tau \vdash q \Rightarrow t : \hat{P}(\tau) \Rightarrow \tau}{\Xi; \Theta \vdash (a, q) \Rightarrow t : (\text{Id} \otimes \hat{P})(\tau) \Rightarrow \tau} \quad \frac{\Xi, a : \tau'; \Theta, a : \tau' \vdash (o, q) \Rightarrow t : (\underline{Q} \otimes \hat{P})(\tau) \Rightarrow \tau}{\Xi; \Theta \vdash (\text{pack}(a, o), q) \Rightarrow t : (\exists a : \tau'. \underline{Q} \otimes \hat{P})(\tau) \Rightarrow \tau}$$

where the right rule simultaneously binds a in both t and Q , and the left rule only binds a in t (but we add a to both contexts to maintain the invariant $\Xi \subseteq \Theta$ for inputs Ξ and Θ). We sort algebra bodies only when a product ends at a unit (possible by design of the functor grammar), and merely under Ξ ; constant functors depend on Θ :

$$\frac{\Xi \vdash t : \tau}{\Xi; \Theta \vdash () \Rightarrow t : I(\tau) \Rightarrow \tau} \quad \frac{\Xi; \Theta \vdash q \Rightarrow t : \hat{P}(\tau) \Rightarrow \tau \quad \Theta \vdash Q \text{ type}[_]}{\Xi; \Theta \vdash (\top, q) \Rightarrow t : (\underline{Q} \otimes \hat{P})(\tau) \Rightarrow \tau}$$

For algebras α of “type” $(F_1 \oplus F_2) \tau \Rightarrow \tau$, we use a straightforward judgment $\alpha \circ \text{inj}_k \stackrel{\triangle}{=} \alpha_k$ (appendix Fig. 5) that outputs the k th clause α_k of input algebra α :

$$\frac{\alpha \circ \text{inj}_1 \stackrel{\triangle}{=} \alpha_1 \quad \alpha \circ \text{inj}_2 \stackrel{\triangle}{=} \alpha_2 \quad \Xi; \Theta \vdash \alpha_1 : F_1(\tau) \Rightarrow \tau \quad \Xi; \Theta \vdash \alpha_2 : F_2(\tau) \Rightarrow \tau}{\Xi; \Theta \vdash \alpha : (F_1 \oplus F_2)(\tau) \Rightarrow \tau}$$

By restricting the bodies of algebras to *index terms* t and the carriers of our F -algebras to *index sorts* τ , we uphold the phase distinction: we can therefore safely refine inductive types by folding them with algebras, and also manage decidable typing.

4.3 Equivalence

We have equivalence judgments for propositions $\Theta \vdash \phi \equiv \psi$ (appendix Fig. 11), logical contexts $\Theta \vdash \Theta_1 \equiv \Theta_2$ (appendix Fig. 12), functors $\Theta \vdash \mathcal{F} \equiv \mathcal{G}$ (appendix Fig. 13), and types $\Theta \vdash A \equiv^\pm B$ (appendix Fig. 14), which use $\Theta \vdash \phi$ true to verify logical equality of index terms. Basically, two entities are equivalent if their respective, structural subparts are equivalent (under the logical context). Type/functor equivalence is used in sum and refined ADT subtyping (type equivalence implies mutual subtyping), as well as to prove algorithmic completeness (appendix Lemma B.108), but context equivalence is only used to prove algorithmic completeness (appendix Lemma B.95).

However, it should be possible to remove equivalence from the system itself, by using “subfunctoring” and covariant sum subtyping. For space reasons, we do not show all their rules here (see appendix), only the ones we think are most likely to surprise.

Refined inductive types are equivalent only if they use syntactically the same algebra (but the algebra must be well-formed at both functors F and G ; this holds by inversion on the conclusion’s presupposed type well-formedness judgments):

$$\frac{\Theta \vdash F \equiv G \quad \Theta \vdash t = t' \text{ true}}{\Theta \vdash \{v : \mu F \mid (\text{fold}_F \alpha) v =_\tau t\} \equiv^+ \{v : \mu G \mid (\text{fold}_G \alpha) v =_\tau t'\}}$$

Two index equality propositions (respectively, two index inequalities) are equivalent if their respective sides are logically equal:

$$\frac{\Theta \vdash t_1 = t'_1 \text{ true} \quad \Theta \vdash t_2 = t'_2 \text{ true}}{\Theta \vdash (t_1 = t_2) \equiv (t'_1 = t'_2)} \quad \frac{\Theta \vdash t_1 = t'_1 \text{ true} \quad \Theta \vdash t_2 = t'_2 \text{ true}}{\Theta \vdash (t_1 \leq t_2) \equiv (t'_1 \leq t'_2)}$$

We use logical context equivalence in proving subsumption admissibility (see Sec. 4.8) and the completeness of algorithmic typing (see Sec. 7.3). Two logical contexts are judgmentally equivalent under Θ if they have exactly the same variable sortings (in the same list positions) and logically equivalent (under Θ) propositions, in the same order. The most interesting rule is the one for propositions, where, in the second premise, we filter out propositions from Θ_1 because we want each respective proposition to be logically equivalent under the propositions (and indexes) of Θ , but variables in Θ_1 (or Θ_2) may appear in ϕ_1 (or ϕ_2):

$$\frac{\Theta \vdash \Theta_1 \equiv \Theta_2 \quad \Theta, \overline{\Theta_1} \vdash \phi_1 \equiv \phi_2}{\Theta \vdash (\Theta_1, \phi_1) \equiv (\Theta_2, \phi_2)}$$

(Note that it is equivalent to use $\overline{\Theta_2}$ rather than $\overline{\Theta_1}$ in the second premise above.)

All equivalence judgments satisfy reflexivity (appendix Lemmas B.75 and B.76), symmetry (appendix Lemmas B.91 and B.92), and transitivity (appendix Lemma B.78).

4.4 Extraction

The judgment $\Theta \vdash A \rightsquigarrow^\pm A' [\Theta_A]$ (Fig. 7) *extracts* quantified variables and \wedge and \supset propositions from the type A , outputting the type A' without these, and the context Θ_A with them. We call A' and Θ_A the type and context *extracted* from A . For negative A , everything is extracted up to an upshift. For positive A , everything is extracted up to any connective that is not \exists , \wedge , or \times . For convenience in program typing (Sec. 4.7), $\Theta \vdash A \rightsquigarrow$ abbreviates $\Theta \vdash A \rightsquigarrow A [\cdot]$ (we sometimes omit the polarity label from extraction judgments). If $\Theta \vdash A \rightsquigarrow$, then we say A is *simple*.

4.5 Subtyping

Declarative subtyping $\Theta \vdash A \leq^\pm B$ is defined in Fig. 8.

Subtyping is polarized into mutually recursive positive $\Theta \vdash P \leq^+ Q$ and negative $\Theta \vdash N \leq^- M$ relations. The design of inference rules for subtyping is guided by sequent calculi, perhaps most clearly seen in the left and right rules pertaining to quantifiers (\exists , \forall), asserting types (\wedge), and guarded types (\supset). This is helpful to establish key properties such as *reflexivity* and *transitivity* (viewing subtyping as a sequent system, we might instead say that the structural *identity* and *cut* rules, respectively, are admissible⁵). We interpret types as sets with some additional structure (Sec. 5), but considering only the sets, we prove that a subtype denotes a subset of the set denoted by any

⁵A proposed inference rule is *admissible* with respect to a system if, whenever the premises of the proposed rule are derivable, we can derive the proposed rule’s conclusion using the system’s inference rules.

$$\boxed{\Theta \vdash A \rightsquigarrow^\pm A' [\Theta_A]} \text{ Under } \Theta, \text{ type } A \text{ extracts to } A' \text{ and } \Theta_A$$

$$\begin{array}{c}
\frac{P \neq \exists, \wedge, \text{ or } \times}{\Theta \vdash P \rightsquigarrow^+ P [\cdot]} \rightsquigarrow^+ \neq \\
\frac{\Theta, a : \tau \vdash P \rightsquigarrow^+ P' [\Theta_P]}{\Theta \vdash \exists a : \tau. P \rightsquigarrow^+ P' [a : \tau, \Theta_P]} \rightsquigarrow^+ \exists \\
\frac{}{\Theta \vdash \uparrow P \rightsquigarrow^- \uparrow P [\cdot]} \rightsquigarrow^- \neq \\
\frac{\Theta, a : \tau \vdash N \rightsquigarrow^- N' [\Theta_N]}{\Theta \vdash \forall a : \tau. N \rightsquigarrow^- N' [a : \tau, \Theta_N]} \rightsquigarrow^- \forall
\end{array}
\quad
\begin{array}{c}
\frac{\Theta \vdash P \rightsquigarrow^+ P' [\Theta_P]}{\Theta \vdash P \wedge \phi \rightsquigarrow^+ P' [\phi, \Theta_P]} \rightsquigarrow^+ \wedge \\
\frac{\Theta \vdash P_1 \rightsquigarrow^+ P'_1 [\Theta_{P_1}] \quad \Theta \vdash P_2 \rightsquigarrow^+ P'_2 [\Theta_{P_2}]}{\Theta \vdash P_1 \times P_2 \rightsquigarrow^+ P'_1 \times P'_2 [\Theta_{P_1}, \Theta_{P_2}]} \rightsquigarrow^+ \times \\
\frac{\Theta \vdash N \rightsquigarrow^- N' [\Theta_N]}{\Theta \vdash \phi \supset N \rightsquigarrow^- N' [\phi, \Theta_N]} \rightsquigarrow^- \supset \\
\frac{\Theta \vdash P \rightsquigarrow^+ P' [\Theta_P] \quad \Theta \vdash N \rightsquigarrow^- N' [\Theta_N]}{\Theta \vdash P \rightarrow N \rightsquigarrow^- P' \rightarrow N' [\Theta_P, \Theta_N]} \rightsquigarrow^- \rightarrow
\end{array}$$

Fig. 7. Declarative extraction

$$\boxed{\Theta \vdash A \leq^\pm B} \text{ Under } \Theta, \text{ type } A \text{ is a subtype of } B$$

$$\begin{array}{c}
\frac{}{\Theta \vdash 1 \leq^+ 1} \leq^+ 1 \qquad \frac{}{\Theta \vdash 0 \leq^+ 0} \leq^+ 0 \\
\frac{\Theta \vdash P_1 \leq^+ Q_1 \quad \Theta \vdash P_2 \leq^+ Q_2}{\Theta \vdash P_1 \times P_2 \leq^+ Q_1 \times Q_2} \leq^+ \times \qquad \frac{\Theta \vdash P_1 \equiv^+ Q_1 \quad \Theta \vdash P_2 \equiv^+ Q_2}{\Theta \vdash P_1 + P_2 \leq^+ Q_1 + Q_2} \leq^+ + \\
\frac{\Theta \vdash P \rightsquigarrow^+ P' [\Theta_P] \quad \Theta_P \neq \cdot}{\Theta \vdash P \leq^+ Q} \leq^+ \rightsquigarrow L \qquad \frac{\Theta, \Theta_P \vdash P' \leq^+ Q}{\Theta \vdash P \leq^+ Q} \leq^+ \rightsquigarrow L \\
\frac{\Theta \vdash P \leq^+ Q \quad \Theta \vdash \phi \text{ true}}{\Theta \vdash P \leq^+ Q \wedge \phi} \leq^+ \wedge R \qquad \frac{\Theta \vdash P \leq^+ [t/a]Q \quad \Theta \vdash t : \tau}{\Theta \vdash P \leq^+ \exists a : \tau. Q} \leq^+ \exists R \\
\frac{\Theta \vdash F \equiv G \quad \Theta \vdash t = t' \text{ true}}{\Theta \vdash \{v : \mu F \mid (\text{fold}_F \alpha) v =_\tau t\} \leq^+ \{v : \mu G \mid (\text{fold}_G \alpha) v =_\tau t'\}} \leq^+ \mu \\
\frac{\Theta \vdash N \leq^- M}{\Theta \vdash \downarrow N \leq^+ \downarrow M} \leq^+ \downarrow \qquad \frac{\Theta \vdash P \leq^+ Q}{\Theta \vdash \uparrow P \leq^- \uparrow Q} \leq^- \uparrow \\
\frac{\Theta \vdash N \leq^- M \quad \Theta \vdash \phi \text{ true}}{\Theta \vdash \phi \supset N \leq^- M} \leq^- \supset L \qquad \frac{\Theta \vdash [t/a]N \leq^- M \quad \Theta \vdash t : \tau}{\Theta \vdash \forall a : \tau. N \leq^- M} \leq^- \forall L \\
\frac{\Theta \vdash M \rightsquigarrow^- M' [\Theta_M] \quad \Theta_M \neq \cdot \quad \Theta, \Theta_M \vdash N \leq^- M'}{\Theta \vdash N \leq^- M} \leq^- \rightsquigarrow R \\
\frac{\Theta \vdash Q \leq^+ P \quad \Theta \vdash N \leq^- M}{\Theta \vdash P \rightarrow N \leq^- Q \rightarrow M} \leq^- \rightarrow
\end{array}$$

Fig. 8. Declarative subtyping

of its supertypes. That is, membership of a (semantic) value in the subtype *implies* its membership in any supertype of the subtype. We may also view subtyping as implication.

Instead of $\leq^+ \rightsquigarrow L$, one might reasonably expect these two rules (the brackets around the rule names indicate that these rules are *not* in our system):

$$\frac{\Theta, \phi \vdash P \leq^+ Q}{\Theta \vdash P \wedge \phi \leq^+ Q} [\leq^+ \wedge L] \qquad \frac{\Theta, a : \tau \vdash P \leq^+ Q}{\Theta \vdash \exists a : \tau. P \leq^+ Q} [\leq^+ \exists L]$$

Similarly, one might expect to have $[\leq^- \supset R]$ and $[\leq^- \forall R]$, dual to the above rules, instead of the dual rule $\leq^- \rightsquigarrow R$. Reading, for example, the above rule $[\leq^+ \wedge L]$ logically and top-down, if Θ and ϕ implies that P implies Q , then we can infer that Θ implies that P and ϕ implies Q . We can also read rules as a bottom-up decision procedure: given $P \wedge \phi$, we know ϕ , so we can assume it; given $\exists a : \tau. P$, we know there exists an index of sort τ such that P , but we don't have a specific index term. However, these rules are not powerful enough to derive reasonable judgments such as $a : \mathbb{N} \vdash 1 \times (1 \wedge a = 3) \leq^+ (1 \wedge a \geq 3) \times 1$: subtyping for the first component requires verifying $a \geq 3$, which is impossible under no logical assumptions. But from a logical perspective, $1 \times (1 \wedge a = 3)$ implies $a \geq 3$. Reading $\leq^+ \rightsquigarrow L$ bottom-up, in this case, we extract $a = 3$ from the subtype, which we later use to verify that $a \geq 3$. The idea is that, for a type in an assumptive position, it does not matter which product component (products are viewed conjunctively) or function argument (in our system, functions must be fully applied to values) to which index data is attached. Moreover, as we'll explain at the end of Sec. 6, the weaker rules by themselves are incompatible with algorithmic completeness. We emphasize that we do *not* include $[\leq^+ \wedge L]$, $[\leq^+ \exists L]$, $[\leq^- \supset R]$ or $[\leq^- \forall R]$ in the system.

For the unit type and the void type, rules $\leq^+ 1$ and $\text{void} \leq^+ 0$ are simply reflexivity. Product subtyping $\leq^+ \times$ is covariant subtyping of component types: a product type is a subtype of another if each component of the former is a subtype of the respective component of the latter. We have covariant shift rules $\leq^+ \downarrow$ and $\leq^- \uparrow$. Function subtyping $\leq^- \rightarrow$ is standard: contravariant (from conclusion to premise, the subtyping direction flips) in the function type's domain and covariant in the function type's codomain.

Rule $\leq^+ \wedge R$ and its dual rule $\leq^- \supset L$ verify the validity of the attached proposition. In rule $\leq^+ \exists R$ and its dual rule $\leq^- \forall L$, we assume that we can conjure a suitable index term t ; in practice (that is, algorithmically), we must introduce an existential variable \hat{a} and then solve it.

Rule $\leq^+ +$ says a sum is a subtype of another sum if their respective subparts are (judgmentally) *equivalent*. Judgmental equivalence does not use judgmental extraction. The logical reading of subtyping begins to clarify why we don't extract anything under a sum connective: $(1 \wedge \text{ff}) + 1$ does not imply ff . However, using equivalence here is a conservative restriction: for example, $(1 \wedge \text{ff}) + (1 \wedge \text{ff})$ does imply ff . Regardless, we don't expect this to be very restrictive in practice because programmers tend not to work with sum types themselves, but rather algebraic inductive types (like μF), and don't need to directly compare, via subtyping, (the unrolling of) different such types (such as the type of lists and the type of natural numbers).

In rule $\leq^+ \mu$, just as in the refined inductive type equivalence rule (Sec. 4.3), a refined inductive type is a subtype of another type if they have judgmentally equivalent functors, they use syntactically the same algebra (that agrees with both subtype and supertype functors), and the index terms on the right-hand side of their measurements are equal under the logical context. As we discuss in Sec. 9, adding polymorphism to the language (future work) might necessitate replacing type and functor equivalence in subtyping with subtyping and "subfunctoring".

In the appendix, we prove that subtyping is reflexive (Lemma B.77) and transitive (Lemma B.83).

Subtyping and type equivalence. We prove that type equivalence implies subtyping (appendix Lemma B.96). To prove that, we use the fact that if Θ_1 is logically equivalent to Θ_2 under their prefix context Θ (judgment $\Theta \vdash \Theta_1 \equiv \Theta_2$) then we can replace Θ_1 with Θ_2 (and vice versa) in derivations (appendix Lemma B.95). We use appendix Lemma B.96 to prove subsumption admissibility (Sec. 4.8) and a subtyping constraint verification transport lemma (mentioned in Sec. 7.2). Conversely, mutual subtyping does not imply type equivalence: $\vdash 1 \wedge \text{tt} \leq 1$ and $\vdash 1 \leq 1 \wedge \text{tt}$ but $\vdash 1 \not\equiv 1 \wedge \text{tt}$ because the unit type is structurally distinct from an asserting type.

4.6 Unrolling

Given $a : \mathbb{N}$, in our system, the type $\text{List } P \ a$ of a -length lists of elements of type P is defined as $\{v : \mu\text{ListF}_P \mid (\text{fold}_{\text{ListF}_P} \text{lenalg}) v = a\}$ where $\text{ListF}_P = I \oplus (P \otimes \text{Id} \otimes I)$ and $\text{lenalg} = \text{inj}_1 () \Rightarrow 0 \mid \text{inj}_2 (\top, (b, ())) \Rightarrow 1 + b$. Assuming we have $\text{succ} : \forall a : \mathbb{N}. \text{Nat}(a) \rightarrow \uparrow \text{Nat}(1 + a)$ for incrementing a (program-level) natural number by one, we define length in our system as follows:

```

rec length : (∀ a : ℕ. List(P)(a) → ↑Nat(a)). λx. match x. {
  into(x') ⇒ match x' {
    inj₁ ⟨⟩ ⇒ -- a = 0
      return into(inj₁ ⟨⟩)
  | inj₂ ⟨⊤, ⟨y, ⟨⟩⟩⟩ ⇒ -- a = 1 + a' such that a' is the length of y
      let z' = length(y);
      let z = succ(z');
      return z
  }
}

```

Checking length against its type annotation, the lambda rule assumes $x : \text{List}(P)(a)$ for an arbitrary $a : \mathbb{N}$. Upon matching x against the pattern $\text{into}(x')$, we know x' should have the *unrolled* type of $\text{List}(P)(a)$. Ignoring refinements, we know that the erasure of this unrolling should be a sum type where the left component represents the empty list and the right component represents a head element together with a tail list. However, in order to verify the refinement that length does what we intend, we need to know more about the length index associated with x —that is, a —in the case where x is nil and in the case where x is a cons cell. Namely, the unrolling of $\text{List}(P)(a)$ should know that $a = 0$ when x is the empty list, and that $a = 1 + a'$ where a' is the length of the tail of x when x is a nonempty list. This is the role of the unrolling judgment, to output just what we need here:

$$\begin{aligned} & \cdot; \vdash \{v : \text{ListF}_P[\mu\text{ListF}_P] \mid \text{lenalg}(\text{ListF}_P(\text{fold}_{\text{ListF}_P} \text{lenalg}) v) =_{\mathbb{N}} a\} \\ & \triangleq (1 \wedge a = 0) + (P \times (\exists a' : \mathbb{N}. \{v : \mu\text{ListF}_P \mid (\text{fold}_{\text{ListF}_P} \text{lenalg}) v =_{\mathbb{N}} a'\} \times (1 \wedge a = 1 + a'))) \end{aligned}$$

That is, the type of P -lists of length a unrolls to either the unit type 1 (representing the empty list) together with the fact that a is 0, or the product of P (the type of the head element) and P -lists (representing the tail) of length a' such that a' is a minus one.

Refined inductive type unrolling $\Xi; \Theta \vdash \{v : G[\mu F] \mid \beta(G(\text{fold}_F \alpha) v) =_{\tau} t\} \triangleq P$, inspired by work in fibrational dependent type theory [Atkey et al. 2012], is defined in Fig. 9. There are two contexts: Ξ is for β and Θ is for G , F , and t . Similarly to algebra well-formedness, we maintain the invariant in unrolling that $\Xi \subseteq \Theta$. The (non-contextual) *input* metavariables are G , F , β , α , τ , and t . The type P , called the *unrolled* type, is an output. As in the list example above, inductive type unrolling is always initiated with $\Xi = \cdot$ and $G = F$ and $\beta = \alpha$.

$\text{Unroll}\oplus$ unrolls each branch and then sums the resulting types. UnrollId outputs the product of the original inductive type but with a measurement given by the recursive result of the fold

$$\boxed{\Xi; \Theta \vdash \{v : G[\mu F] \mid \beta(G(\text{fold}_F \alpha) v) =_\tau t\} \doteq P} \quad \text{Abbreviated}$$

$$\Xi; \Theta \vdash \text{unroll}_{F,\alpha}(G; \beta; \tau; t) \doteq P$$

$$\frac{\beta \circ \text{inj}_1 \doteq \beta_1 \quad \Xi; \Theta \vdash \{v : G[\mu F] \mid \beta_1(G(\text{fold}_F \alpha) v) =_\tau t\} \doteq P \quad \beta \circ \text{inj}_2 \doteq \beta_2 \quad \Xi; \Theta \vdash \{v : H[\mu F] \mid \beta_2(H(\text{fold}_F \alpha) v) =_\tau t\} \doteq Q}{\Xi; \Theta \vdash \{v : (G \oplus H)[\mu F] \mid \beta((G \oplus H)(\text{fold}_F \alpha) v) =_\tau t\} \doteq P + Q} \text{Unroll}\oplus$$

$$\frac{\Xi, a : \tau; \Theta, a : \tau \vdash \{v : \hat{P}[\mu F] \mid (q \Rightarrow t') (\hat{P}(\text{fold}_F \alpha) v) =_\tau t\} \doteq Q}{\Xi; \Theta \vdash \left\{ v : (\text{Id} \otimes \hat{P})[\mu F] \mid ((a, q) \Rightarrow t') ((\text{Id} \otimes \hat{P})(\text{fold}_F \alpha) v) =_\tau t \right\} \doteq Q} \text{Unroll}\text{Id}$$

$$\Xi; \Theta \vdash \left\{ v : (\text{Id} \otimes \hat{P})[\mu F] \mid ((a, q) \Rightarrow t') ((\text{Id} \otimes \hat{P})(\text{fold}_F \alpha) v) =_\tau t \right\} \doteq \exists a : \tau. \{v : \mu F \mid (\text{fold}_F \alpha) v =_\tau a\} \times Q$$

$$\frac{\Xi, a : \tau'; \Theta, a : \tau' \vdash \{v : (\underline{Q} \otimes \hat{P})[\mu F] \mid ((o, q) \Rightarrow t') ((\underline{Q} \otimes \hat{P})(\text{fold}_F \alpha) v) =_\tau t\} \doteq Q'}{\Xi; \Theta \vdash \left\{ v : (\exists a : \tau'. \underline{Q} \otimes \hat{P})[\mu F] \mid ((\text{pack}(a, o), q) \Rightarrow t') ((\exists a : \tau'. \underline{Q} \otimes \hat{P})(\text{fold}_F \alpha) v) =_\tau t \right\} \doteq Q'} \text{Unroll}\exists$$

$$\Xi; \Theta \vdash \left\{ v : (\exists a : \tau'. \underline{Q} \otimes \hat{P})[\mu F] \mid ((\text{pack}(a, o), q) \Rightarrow t') ((\exists a : \tau'. \underline{Q} \otimes \hat{P})(\text{fold}_F \alpha) v) =_\tau t \right\} \doteq \exists a : \tau'. Q'$$

$$\frac{\Xi; \Theta \vdash \{v : \hat{P}[\mu F] \mid (q \Rightarrow t') (\hat{P}(\text{fold}_F \alpha) v) =_\tau t\} \doteq Q'}{\Xi; \Theta \vdash \{v : (\underline{Q} \otimes \hat{P})[\mu F] \mid ((\top, q) \Rightarrow t') ((\underline{Q} \otimes \hat{P})(\text{fold}_F \alpha) v) =_\tau t\} \doteq Q \times Q'} \text{Unroll}\text{Const}$$

$$\frac{}{\Xi; \Theta \vdash \{v : I[\mu F] \mid ((\ () \Rightarrow t') (I(\text{fold}_F \alpha) v) =_\tau t\} \doteq 1 \wedge (t = t')} \text{Unroll}\text{I}$$

Fig. 9. Unrolling

(over which we existentially quantify), together with the rest of the unrolling. $\text{Unroll}\exists$ pushes the packed index variable a onto the context and continues unrolling, existentially quantifying over the result; in the conclusion, a is simultaneously bound in Q and t' . $\text{Unroll}\text{Const}$ outputs a product of the type of the constant functor and the rest of the unrolling. UnrollI simply outputs the unit type together with the index term equality given by the (unrolled) measurement.

If our functor and algebra grammars were instead more direct, like those implicitly used in the introduction (Sec. 1) and overview (Sec. 2), and explicitly discussed in Sec. 4.0.1, then we would have to modify the unrolling judgment, and it would need two more rules. We expect everything would still work, but we prefer having to consider fewer rules when proving metatheory.

Unrolling, equivalence and subtyping. Substituting judgmentally equivalent types, functors and indexes for the inputs of unrolling generates an output type that is both a subtype and supertype of the original unrolling output:

LEMMA 4.1 (UNROLL TO MUTUAL SUBTYPES).

(Lemma B.97 in appendix)

If $\Xi; \Theta \vdash \{v : G[\mu F] \mid \beta(G(\text{fold}_F \alpha) v) =_\tau t\} \doteq P$

and $\Theta \vdash G \equiv G'$ and $\Theta \vdash F \equiv F'$ and $\Theta \vdash t = t'$ true,

then there exists Q such that $\Xi; \Theta \vdash \{v : G'[\mu F'] \mid \beta(G'(\text{fold}_{F'} \alpha) v) =_\tau t'\} \doteq Q$

and $\Theta \vdash P \leq^+ Q$ and $\Theta \vdash Q \leq^+ P$.

We use this to prove subsumption admissibility (see Sec. 4.8) for the cases that involve constructing and pattern matching inductive values.

4.7 Typing

Declarative bidirectional typing rules are given in Figures 10, 11, and 12. By careful design, guided by logical principles, all typing rules are syntax-directed. That is, when deriving a conclusion, at most one rule is compatible with the syntax of the input program term and the principal input type.

To manage the interaction between subtyping and program typing, types in a well-formed (under Θ) program context Γ must be invariant under extraction: for all $(x : P) \in \Gamma$, we have $\Theta \vdash P \rightsquigarrow^+ P [\cdot]$ (that is, $\Theta \vdash P \rightsquigarrow$). We maintain this invariant in program typing by extracting before adding any variable typings to the context.

$$\boxed{\Theta; \Gamma \vdash h \Rightarrow P} \quad \text{Under } \Theta \text{ and } \Gamma, \text{ head } h \text{ synthesizes type } P$$

$$\frac{(x : P) \in \Gamma}{\Theta; \Gamma \vdash x \Rightarrow P} \text{Decl} \Rightarrow \text{Var} \qquad \frac{\Theta \vdash P \text{ type}[\Xi] \quad \Theta; \Gamma \vdash v \Leftarrow P}{\Theta; \Gamma \vdash (v : P) \Rightarrow P} \text{Decl} \Rightarrow \text{ValAnnot}$$

$$\boxed{\Theta; \Gamma \vdash g \Rightarrow \uparrow P} \quad \text{Under } \Theta \text{ and } \Gamma, \text{ bound expression } g \text{ synthesizes type } \uparrow P$$

$$\frac{\Theta; \Gamma \vdash h \Rightarrow \downarrow N \quad \Theta; \Gamma; [N] \vdash s \gg \uparrow P}{\Theta; \Gamma \vdash h(s) \Rightarrow \uparrow P} \text{Decl} \Rightarrow \text{App}$$

$$\frac{\Theta \vdash P \text{ type}[\Xi] \quad \Theta; \Gamma \vdash e \Leftarrow \uparrow P}{\Theta; \Gamma \vdash (e : \uparrow P) \Rightarrow \uparrow P} \text{Decl} \Rightarrow \text{ExpAnnot}$$

Fig. 10. Declarative head and bound expression type synthesis

The judgment $\Theta; \Gamma \vdash h \Rightarrow P$ (Fig. 10) synthesizes the type P from the head h . This judgment is synthesizing, because it is used in what are, from a Curry–Howard perspective, kinds of cut rules: $\text{Decl} \Rightarrow \text{App}$ and $\text{Decl} \Leftarrow \text{match}$, discussed later. The synthesized type is the cut type, which does not appear in the conclusion of $\text{Decl} \Rightarrow \text{App}$ or $\text{Decl} \Leftarrow \text{match}$. For head variables, we look up the variable’s type in the context Γ ($\text{Decl} \Rightarrow \text{Var}$). For annotated values, we synthesize the annotation ($\text{Decl} \Rightarrow \text{ValAnnot}$).

The judgment $\Theta; \Gamma \vdash g \Rightarrow \uparrow P$ (Fig. 10) synthesizes the type $\uparrow P$ from the bound expression g . Similarly to the synthesizing judgment for heads, this judgment is synthesizing because it is used in a cut rule $\text{Decl} \Leftarrow \text{let}$ (the synthesized type is again the cut type). Bound expressions only synthesize an upshift because of their (lone) role in rule $\text{Decl} \Leftarrow \text{let}$, discussed later. For an application of a head to a spine ($\text{Decl} \Rightarrow \text{App}$, an auxiliary cut rule), we first synthesize the head’s type (which must be a downshift), and then check the spine against the thunked computation type, synthesizing the latter’s return type. (Function applications must always be fully applied, but we can simulate partial application via η -expansion. For example, given $x : P_1$ and $h \Rightarrow \downarrow (P_1 \rightarrow P_2 \rightarrow \uparrow Q)$, to partially apply h to x we can write $\lambda y. \text{let } z = h(x, y); \dots$) For annotated expressions, we synthesize the annotation ($\text{Decl} \Rightarrow \text{ExpAnnot}$), which must be an upshift. If the programmer wants, say, to verify guard constraints in N of an expression e of type N whenever it is run, then they must annotate it: $(\text{return } \{e\} : \uparrow \downarrow N)$. If an e of type N is intended to be a function to be applied (as a head to a spine; $\text{Decl} \Rightarrow \text{App}$) only if the guards of N can be verified and the universally quantified indexes of N can be instantiated, then the programmer must thunk and annotate it: $(\{e\} : \downarrow N)$. The two annotation rules have explicit type well-formedness premises to emphasize that type annotations are provided by the programmer.

The judgment $\Theta; \Gamma \vdash v \Leftarrow P$ (Fig. 11) checks the value v against the type P . From a Curry–Howard perspective, this judgment corresponds to a *right-focusing* stage. According to rule $\text{Decl} \Leftarrow \exists$, a value

$\Theta; \Gamma \vdash v \Leftarrow P$

Under Θ and Γ , value v checks against type P

$$\frac{P \neq \exists, \wedge \quad (x : Q) \in \Gamma \quad \Theta \vdash Q \leq^+ P}{\Theta; \Gamma \vdash x \Leftarrow P} \text{Decl} \Leftarrow \text{Var} \qquad \frac{}{\Theta; \Gamma \vdash \langle \rangle \Leftarrow 1} \text{Decl} \Leftarrow 1$$

$$\frac{\Theta; \Gamma \vdash v_1 \Leftarrow P_1 \quad \Theta; \Gamma \vdash v_2 \Leftarrow P_2}{\Theta; \Gamma \vdash \langle v_1, v_2 \rangle \Leftarrow P_1 \times P_2} \text{Decl} \Leftarrow \times \qquad \frac{\Theta; \Gamma \vdash v \Leftarrow P_k}{\Theta; \Gamma \vdash \text{inj}_k v \Leftarrow P_1 + P_2} \text{Decl} \Leftarrow +_k$$

$$\frac{\Theta; \Gamma \vdash v \Leftarrow [t/a]P \quad \Theta \vdash t : \tau}{\Theta; \Gamma \vdash v \Leftarrow (\exists a : \tau. P)} \text{Decl} \Leftarrow \exists \qquad \frac{\Theta; \Gamma \vdash v \Leftarrow P \quad \Theta \vdash \phi \text{ true}}{\Theta; \Gamma \vdash v \Leftarrow P \wedge \phi} \text{Decl} \Leftarrow \wedge$$

$$\frac{; \Theta \vdash \{v : F[\mu F] \mid \alpha (F (\text{fold}_F \alpha) v) =_\tau t\} \doteq P \quad \Theta; \Gamma \vdash v \Leftarrow P}{\Theta; \Gamma \vdash \text{into}(v) \Leftarrow \{v : \mu F \mid (\text{fold}_F \alpha) v =_\tau t\}} \text{Decl} \Leftarrow \mu$$

$$\frac{\Theta; \Gamma \vdash e \Leftarrow N}{\Theta; \Gamma \vdash \{e\} \Leftarrow \downarrow N} \text{Decl} \Leftarrow \downarrow$$

$\Theta; \Gamma \vdash e \Leftarrow N$

Under Θ and Γ , expression e checks against type N

$$\frac{\Theta; \Gamma \vdash v \Leftarrow P}{\Theta; \Gamma \vdash \text{return } v \Leftarrow \uparrow P} \text{Decl} \Leftarrow \uparrow$$

$$\frac{\Theta \vdash N \not\rightsquigarrow \quad \Theta; \Gamma \vdash g \Rightarrow \uparrow P \quad \Theta \vdash P \rightsquigarrow^+ P' [\Theta_P] \quad \Theta, \Theta_P; \Gamma, x : P' \vdash e \Leftarrow N}{\Theta; \Gamma \vdash \text{let } x = g; e \Leftarrow N} \text{Decl} \Leftarrow \text{let}$$

$$\frac{\Theta \vdash N \not\rightsquigarrow \quad \Theta; \Gamma \vdash h \Rightarrow P \quad \Theta; \Gamma; [P] \vdash \{r_i \Rightarrow e_i\}_{i \in I} \Leftarrow N}{\Theta; \Gamma \vdash \text{match } h \{r_i \Rightarrow e_i\}_{i \in I} \Leftarrow N} \text{Decl} \Leftarrow \text{match}$$

$$\frac{\Theta \vdash P \rightarrow N \not\rightsquigarrow \quad \Theta; \Gamma, x : P \vdash e \Leftarrow N}{\Theta; \Gamma \vdash \lambda x. e \Leftarrow P \rightarrow N} \text{Decl} \Leftarrow \lambda$$

$$\frac{\Theta \vdash N \not\rightsquigarrow \quad \Theta \vdash \text{ff true}}{\Theta; \Gamma \vdash \text{unreachable} \Leftarrow N} \text{Decl} \Leftarrow \text{Unreachable}$$

$$\frac{\Theta \vdash N \not\rightsquigarrow \quad \Theta \vdash \forall a : \mathbb{N}. M \leq^- N \quad \Theta, a : \mathbb{N}; \Gamma, x : \downarrow (\forall a' : \mathbb{N}. (a' < a) \supset [a'/a]M) \vdash e \Leftarrow M}{\Theta; \Gamma \vdash \text{rec } x : (\forall a : \mathbb{N}. M). e \Leftarrow N} \text{Decl} \Leftarrow \text{rec}$$

$$\frac{\Theta \vdash N \rightsquigarrow N' [\Theta_N] \quad \Theta_N \neq \cdot \quad \Theta, \Theta_N; \Gamma \vdash e \Leftarrow N'}{\Theta; \Gamma \vdash e \Leftarrow N} \text{Decl} \Leftarrow \rightsquigarrow$$

Fig. 11. Declarative value and expression type checking

checks against an existential type if there is an index instantiation it checks against (declaratively, an index is conjured, but algorithmically we will have to solve for one). For example, as discussed in Sec. 4.2, checking the program value one representing 1 against type $\exists a : \mathbb{N}. \text{Nat}(a)$ solves a to an index semantically equal to 1. According to rule $\text{Decl} \Leftarrow \wedge$, a value checks against an asserting type if the asserted proposition ϕ holds (and the value checks against the type to which ϕ is

$$\boxed{\Theta; \Gamma; [P] \vdash \{r_i \Rightarrow e_i\}_{i \in I} \Leftarrow N} \quad \text{Under } \Theta \text{ and } \Gamma, \text{ patterns } r_i \text{ match against (input) type } P \text{ and branch expressions } e_i \text{ check against type } N$$

$$\frac{\Theta, a : \tau; \Gamma; [P] \vdash \{r_i \Rightarrow e_i\}_{i \in I} \Leftarrow N}{\Theta; \Gamma; [\exists a : \tau. P] \vdash \{r_i \Rightarrow e_i\}_{i \in I} \Leftarrow N} \text{DeclMatch}\exists$$

$$\frac{\Theta, \phi; \Gamma; [P] \vdash \{r_i \Rightarrow e_i\}_{i \in I} \Leftarrow N}{\Theta; \Gamma; [P \wedge \phi] \vdash \{r_i \Rightarrow e_i\}_{i \in I} \Leftarrow N} \text{DeclMatch}\wedge \qquad \frac{\Theta; \Gamma \vdash e \Leftarrow N}{\Theta; \Gamma; [1] \vdash \{\langle \rangle \Rightarrow e\} \Leftarrow N} \text{DeclMatch}1$$

$$\frac{\Theta \vdash P_1 \rightsquigarrow P'_1 [\Theta_1] \quad \Theta \vdash P_2 \rightsquigarrow P'_2 [\Theta_2] \quad \Theta, \Theta_1, \Theta_2; \Gamma, x_1 : P'_1, x_2 : P'_2 \vdash e \Leftarrow N}{\Theta; \Gamma; [P_1 \times P_2] \vdash \{\langle x_1, x_2 \rangle \Rightarrow e\} \Leftarrow N} \text{DeclMatch}\times$$

$$\frac{\Theta \vdash P_1 \rightsquigarrow P'_1 [\Theta_1] \quad \Theta, \Theta_1; \Gamma, x_1 : P'_1 \vdash e_1 \Leftarrow N \quad \Theta \vdash P_2 \rightsquigarrow P'_2 [\Theta_2] \quad \Theta, \Theta_2; \Gamma, x_2 : P'_2 \vdash e_2 \Leftarrow N}{\Theta; \Gamma; [P_1 + P_2] \vdash \{\text{inj}_1 x_1 \Rightarrow e_1 \mid \text{inj}_2 x_2 \Rightarrow e_2\} \Leftarrow N} \text{DeclMatch}+$$

$$\frac{}{\Theta; \Gamma; [0] \vdash \{\} \Leftarrow N} \text{DeclMatch}0$$

$$\frac{; \Theta \vdash \{v : F[\mu F] \mid \alpha (F (\text{fold}_F \alpha) v) =_\tau t\} \doteq Q \quad \Theta \vdash Q \rightsquigarrow Q' [\Theta_Q] \quad \Theta, \Theta_Q; \Gamma, x : Q' \vdash e \Leftarrow N}{\Theta; \Gamma; [\{v : \mu F \mid (\text{fold}_F \alpha) v =_\tau t\}] \vdash \{\text{into}(x) \Rightarrow e\} \Leftarrow N} \text{DeclMatch}\mu$$

$$\boxed{\Theta; \Gamma; [N] \vdash s \gg \uparrow P} \quad \text{Under } \Theta \text{ and } \Gamma, \text{ if a head of type } \downarrow N \text{ is applied to the spine } s, \text{ then it will return a result of type } \uparrow P$$

$$\frac{\Theta \vdash t : \tau \quad \Theta; \Gamma; [[t/a]N] \vdash s \gg \uparrow P}{\Theta; \Gamma; [\forall a : \tau. N] \vdash s \gg \uparrow P} \text{DeclSpine}\forall$$

$$\frac{\Theta \vdash \phi \text{ true} \quad \Theta; \Gamma; [N] \vdash s \gg \uparrow P}{\Theta; \Gamma; [\phi \supset N] \vdash s \gg \uparrow P} \text{DeclSpine}\supset$$

$$\frac{\Theta; \Gamma \vdash v \Leftarrow Q \quad \Theta; \Gamma; [N] \vdash s \gg \uparrow P}{\Theta; \Gamma; [Q \rightarrow N] \vdash v, s \gg \uparrow P} \text{DeclSpine}App \qquad \frac{}{\Theta; \Gamma; [\uparrow P] \vdash \cdot \gg \uparrow P} \text{DeclSpine}Nil$$

Fig. 12. Declarative pattern matching and spine typing

connected). Instead of a general value type subsumption rule like

$$\frac{\Theta; \Gamma \vdash v \Leftarrow Q \quad \Theta \vdash Q \leq^+ P}{\Theta; \Gamma \vdash v \Leftarrow P}$$

we restrict subsumption to (value) variables, and prove that subsumption is admissible (see Section 4.8). This is easier to implement efficiently because the type checker would otherwise have to guess Q (and possibly need to backtrack), whereas $\text{Decl}\Leftarrow\text{Var}$ need only look up the variable. Further, the $P \neq \exists, \wedge$ constraint on $\text{Decl}\Leftarrow\text{Var}$ means that any top-level \exists or \wedge constraints must be verified before subtyping, eliminating nondeterminism of verifying these in subtyping or typing. Rule $\text{Decl}\Leftarrow\mu$ checks the unrolled value against the unrolled inductive type. Rule $\text{Decl}\Leftarrow 1$ says $\langle \rangle$ checks against 1. Rule $\text{Decl}\Leftarrow\times$ says a pair checks against a product if each pair component checks

against its corresponding factor. Rule $\text{Decl} \Leftarrow +_k$ says a value injected into the k th position checks against a sum if it can be checked against the k th summand. Rule $\text{Decl} \Leftarrow \downarrow$ checks the thunked expression against the computation type N under the given thunk type $\downarrow N$.

The judgment $\Theta; \Gamma \vdash e \Leftarrow N$ (Fig. 11) checks the expression e against the type N . From a Curry–Howard perspective, this judgment is a *right-inversion* stage with *stable* moments ($\text{Decl} \Leftarrow \text{let}$ and $\text{Decl} \Leftarrow \text{match}$, which enter left- or right-focusing stages, respectively). Instead of $\text{Decl} \Leftarrow \rightsquigarrow$, one might expect two rules (one for \forall and one for \supset) that simply put the universal index variable or proposition into logical context, but these alone are less compatible with subsumption admissibility (see Sec. 4.8) due to the use of extraction in subtyping rule $\leq^- \rightsquigarrow R$. However, the idea is still the same: here we are using indexes, not verifying them as in the dual left-focusing stage. To reduce $\text{Decl} \Leftarrow \rightsquigarrow$ nondeterminism, and to enable a formal correspondence between our system and (a variant of) CBPV (which has a general \downarrow elimination rule), the other (expression) rules must check against a *simple* type. In practice, we eagerly apply (if possible) $\text{Decl} \Leftarrow \rightsquigarrow$ immediately when type checking an expression; extracted types are invariant under extraction.

All applications $h(s)$ must be named and sequenced via $\text{Decl} \Leftarrow \text{let}$, which we may think of as monadic binding, and is a key cut rule. Other computations—annotated returner expressions ($e : \uparrow P$)—must also be named and sequenced via $\text{Decl} \Leftarrow \text{let}$. It would not make sense to allow arbitrary negative annotations because that would require verifying constraints and instantiating indexes that should only be done when the annotated expression is applied, which does not occur in $\text{Decl} \Leftarrow \text{let}$ itself.

Heads, that is, head variables and annotated values, can be pattern matched via $\text{Decl} \Leftarrow \text{match}$. From a Curry–Howard perspective, the rule $\text{Decl} \Leftarrow \text{match}$ is a cut rule dual to the cut rule $\text{Decl} \Leftarrow \text{let}$: the latter binds the result of a computation to a (sequenced) computation, whereas the former binds the deconstruction of a value to, and directs control flow of, a computation. Rule $\text{Decl} \Leftarrow \lambda$ is standard (besides the check that $P \rightarrow N$ is simple). Rule $\text{Decl} \Leftarrow \text{rec}$ requires an annotation that universally quantifies over the argument a that must be smaller at each recursive call, as dictated by its annotation in the last premise: $x : \downarrow (\forall a' : \mathbb{N}. (a' < a) \supset [a'/a]M)$ only allows x to be used for $a' < a$, ensuring that *refined* recursive functions are well-founded (according to $<$ on naturals). Rule $\text{Decl} \Leftarrow \uparrow$ checks that the value being returned has the positive type under the given returner type (\uparrow); this may be thought of as a monadic return operation. Rule $\text{Decl} \Leftarrow \text{Unreachable}$ says that unreachable checks against any type, provided the logical context is inconsistent; for example, an impossible pattern in pattern matching extracts to an inconsistent context.

Rule $\text{Decl} \Leftarrow \text{rec}$ only handles one termination metric, namely $<$ on natural numbers. This is only to simplify our presentation, and is not a fundamental limitation of the system. We can, for example, add a rule that encodes a termination metric $<$ on the sum of two natural numbers:

$$\frac{\Theta \vdash N \rightsquigarrow \quad \Theta \vdash \forall a : \mathbb{N}. \forall b : \mathbb{N}. M \leq^- N \quad \Theta, a : \mathbb{N}, b : \mathbb{N}; \Gamma, x : \downarrow (\forall a' : \mathbb{N}. \forall b' : \mathbb{N}. (a' + b' < a + b) \supset [a'/a][b'/b]M) \vdash e \Leftarrow M}{\Theta; \Gamma \vdash \text{rec } x : (\forall a : \mathbb{N}. \forall b : \mathbb{N}. M). e \Leftarrow N}$$

It is somewhat straightforward to update the metatheory for the system with this rule added. This rule obviates, for example, the ghost parameter used in the mergesort example of Sec. 3. Similarly, one could add rules for other termination metrics, such as lexicographic induction.

The judgment $\Theta; \Gamma; [P] \vdash \{r_i \Rightarrow e_i\}_{i \in I} \Leftarrow N$ (Fig. 12) decomposes P , according to patterns r_i (if $P \neq \wedge$ or \exists , which have no computational content; if $P = \wedge$ or \exists , the index is put in logical context for use), and checks that each branch e_i has type N . The rules are straightforward. Indexes from matching on existential and asserting types are used, not verified (as in value typechecking); we deconstruct heads, and to synthesize a type for a head, its indexes must hold, so within the

pattern matching stage itself, we may assume and use them. From a Curry–Howard perspective, this judgment corresponds to a *left-inversion* stage. However, it is not *strongly* focused, that is, it does not decompose P eagerly and as far as possible; therefore, “stage” might be slightly misleading. If our system were more strongly focused, we would have nested patterns, at least for all positive types except inductive types; it’s unclear how strong focusing on inductive types would work.

The judgment $\Theta; \Gamma; [N] \vdash s \gg \uparrow P$ (Fig. 12) checks the spine s against N , synthesizing the return type $\uparrow P$. From a Curry–Howard perspective, this judgment corresponds to a *left-focusing* stage. The rules are straightforward: decompose the given N , checking index constraints (DeclSpine \forall and DeclSpine \supset) and values (DeclSpineApp) until an upshift, the return type, is synthesized (DeclSpineNil). Similarly to dual rule Decl $\Leftarrow\exists$, the declarative rule DeclSpine \forall conjures an index measuring a value, but in this case an argument value in a spine. For example, in applying a head of type $\forall a : \mathbb{N}. \text{Nat}(a) \rightarrow \uparrow \text{Nat}(a)$ to the spine with program value one representing 1, we must instantiate a to an index semantically equal to 1; we show how this works algorithmically in Sec. 6.4. All universal quantifiers (in the input type of a spine judgment) are solvable algorithmically, because in a well-formed return type, the set of value-determined indexes Ξ is empty.

4.8 Substitution

A key correctness result that we prove is a substitution lemma: substitution (of index terms for index variables and program values for program variables) preserves typing. We now extend the index-level syntactic substitutions (and the sequential substitution operation) introduced in Sec. 4.1. A *syntactic substitution* $\sigma ::= \cdot \mid \sigma, t/a \mid \sigma, v : P/x$ is essentially a list of terms to be substituted for variables. Substitution application $[\sigma]-$ is a sequential substitution metaoperation on types and terms. On program terms, it is a kind of *hereditary substitution*⁶ [Watkins et al. 2004; Pfenning 2008] in the sense that, at head variables (note the h superscript in the Fig. 13 definition; we elide h when clear from context), an annotation is produced if the value and the head variable being replaced by it are not equal—thereby modifying the syntax tree of the substitutee. Otherwise, substitution is standard (homomorphic application) and does not use the value’s associated type given in σ : see Fig. 13.

In the definition given in Fig. 13, an annotation is not produced if $v = x$ so that $x : P/x$ is always an *identity* substitution: that is, $[x : P/x]^h x = x$. As usual, we assume variables are α -renamed to avoid capture by substitution.

The judgment $\Theta_0; \Gamma_0 \vdash \sigma : \Theta; \Gamma$ (appendix Fig. 9) means that, under Θ_0 and Γ_0 , we know σ is a substitution of index terms and program values for variables in Θ and Γ , respectively. The key rule of this judgment is for program value entries (the three elided rules are similar to the three rules for syntactic substitution typing at index level, found near the start of Sec. 4.1, but adds program contexts Γ where appropriate):

$$\frac{\Theta_0; \Gamma_0 \vdash \sigma : \Theta; \Gamma \quad \Theta_0; \Gamma_0 \vdash [\sigma]v \Leftarrow [[\sigma]]P \quad x \notin \text{dom}(\Gamma)}{\Theta_0; \Gamma_0 \vdash (\sigma, v : P/x) : \Theta; \Gamma, x : P}$$

We apply the rest of the syntactic substitution—that is, the σ in the rule—to v and P because the substitution operation is sequential; v may mention variables in Γ and Θ , and P may mention variables in Θ . The metaoperation $[-]$ filters out program variable entries (program variables cannot appear in types, functors, algebras or indexes).

That substitution respects typing is an important correctness property of the type system. We state only two parts here, but those of the remaining program typing judgments are similar; all six parts are mutually recursive.

⁶Typically, hereditary substitution reduces terms after substitution, modifying the syntax tree.

$$\begin{aligned}
[v : P/x]^h y &= y \quad (\text{if } y \neq x) \\
[v : P/x]^h x &= \begin{cases} x & \text{if } v = x \\ (v : P) & \text{else} \end{cases} \\
[v : P/x]^h (v_0 : P_0) &= ([v : P/x]v_0 : P_0) \\
[v : P/x](h(s)) &= ([v : P/x]^h h)([v : P/x]s) \\
[v : P/x](e : \uparrow Q) &= ([v : P/x]e : \uparrow Q) \\
[v : P/x]y &= y \quad (\text{if } y \neq x) \\
[v : P/x]x &= v \\
[v : P/x]\langle v_1, v_2 \rangle &= \langle [v : P/x]v_1, [v : P/x]v_2 \rangle \\
&\vdots \\
[v : P/x](\text{match } h \{r_i \Rightarrow e_i\}_{i \in I}) &= \text{match } ([v : P/x]^h h) ([v : P/x]\{r_i \Rightarrow e_i\}_{i \in I}) \\
&\vdots
\end{aligned}$$

Fig. 13. Definition of syntactic substitution on program terms

LEMMA 4.2 (SYNTACTIC SUBSTITUTION).

*(Lemma B.107 in appendix)*Assume $\Theta_0; \Gamma_0 \vdash \sigma : \Theta; \Gamma$.

- (1) If $\Theta; \Gamma \vdash h \Rightarrow P$, then there exists Q such that $\Theta_0 \vdash Q \leq^+ [[\sigma]]P$ and $\Theta_0; \Gamma_0 \vdash [\sigma]h \Rightarrow Q$.
- (2) If $\Theta; \Gamma \vdash e \Leftarrow N$, then $\Theta_0; \Gamma_0 \vdash [\sigma]e \Leftarrow [[\sigma]]N$.

In part (1), if substitution creates a head variable with stronger type, then the stronger type Q is synthesized. The proof relies on other structural properties such as weakening. It also relies on subsumption admissibility, which captures what we mean by “stronger type”. We show only one part; the mutually recursive parts for the other five program typing judgments are similar.

LEMMA 4.3 (SUBSUMPTION ADMISSIBILITY).

*(Lemma B.106 in appendix)*Assume $\Theta \vdash \Gamma' \leq \Gamma$ (pointwise subtyping).

- (1) If $\Theta; \Gamma \vdash v \Leftarrow P$ and $\Theta \vdash P \leq^+ Q$, then $\Theta; \Gamma' \vdash v \Leftarrow Q$.

Subtypes are stronger than supertypes. That is, if we can check a value against a type, then we know that it also checks against any of the type’s supertypes; similarly for expressions. Pattern matching is similar, but it also says we can match on a stronger type. A head or bound expression can synthesize a stronger type under a stronger context. Similarly, with a stronger input type, a spine can synthesize a stronger return type.

5 TYPE SOUNDNESS

We prove type (and substitution) soundness of the declarative system with respect to an elementary domain-theoretic denotational semantics. Refined type soundness implies the refined system’s totality and logical consistency.

Refinement type systems refine already-given type systems, and the soundness of the former depends on that of the latter [Melliès and Zeilberger 2015]. Thus, the semantics of our refined

system is defined in terms of that of its underlying, unrefined system, which we discuss in Section 5.1.

Notation: We define the disjoint union $X \uplus Y$ of sets X and Y by $X \uplus Y = (\{1\} \times X) \cup (\{2\} \times Y)$ and define $inj_k : X_k \rightarrow X_1 \uplus X_2$ by $inj_k(d) = (k, d)$. Semantic values are usually named d, f, g , or V .

5.1 Unrefined System

For space reasons, we do not fully present the unrefined system and its semantics here (see appendix Sec. A.4). The unrefined system is basically just the refined system with everything pertaining to indexes erased. The program terms of the unrefined system have almost the same syntax as those of the refined system, but an unrefined, recursive expression has no type annotation, and we replace the expression `unreachable` by `diverge`, which stands for an inexhaustive pattern-matching error. The unrefined system satisfies a substitution lemma (appendix Lemma C.1) similar to that of the refined system, but its proof is simpler and does not rely on subsumption admissibility, because the unrefined system has no subtyping.

In CBPV, nontermination is regarded as an effect, so value and computation types denote different kinds of mathematical things: predomains and domains, respectively [Levy 2004], which are both sets with some structure. Because we have recursive expressions, we must model nontermination, an effect. We use elementary domain theory. For our (unrefined) system, we interpret (unrefined) positive types as predomains and (unrefined) negative types as domains. The only effect we consider in this paper is nontermination (though we simulate inexhaustive pattern-matching errors with it); we take (chain-)complete partial orders (cpo) as predomains, and pointed (chain-)complete partial orders (cppo) as domains.

Positive types and functors. The grammar for unrefined positive types is similar to that for refined positive types, but lacks asserting and existential types, and unrefined inductive types μF are not refined by predicates. Unrefined inductive types use the unrefined functor grammar, which is the same as the refined functor grammar but uses unrefined types in constant functors.

$$P, Q ::= 1 \mid P \times Q \mid 0 \mid P + Q \mid \downarrow N \mid \mu F$$

The denotations of unrefined positive types are standard. We briefly describe their partial orders, then describe the meaning of functors, and lastly return to the meaning of inductive types (which involve functors).

We give (the denotation of) 1 (denoting the distinguished terminal object $\{\bullet\}$) the discrete order $\{(\bullet, \bullet)\}$. For $P \times Q$ (denoting product) we use component-wise order $((d_1, d_2) \sqsubseteq_{D_1 \times D_2} (d'_1, d'_2)$ if $d_1 \sqsubseteq_{D_1} d'_1$ and $d_2 \sqsubseteq_{D_2} d'_2$), for 0 (denoting the initial object) we use the empty order, and for $P + Q$ (denoting coproduct, that is, disjoint union \uplus) we use injection-wise order $(inj_j d \sqsubseteq_{D_1 \uplus D_2} inj_k d'$ if $j = k$ and $d \sqsubseteq_{D_j} d'$). We give $\downarrow N$ the order of N , that is, \downarrow denotes the forgetful functor from the category **Cppo** of cppos and continuous functions to the category **Cpo** of cpos and continuous functions. Finally, $V_1 \sqsubseteq_{\llbracket \mu F \rrbracket} V_2$ if $V_1 \sqsubseteq_{\llbracket F \rrbracket^{k+1}\theta} V_2$ for some $k \in \mathbb{N}$, inheriting the type denotation orders as the functor is applied.

The denotations of unrefined functors are standard **Cpo** endofunctors. We briefly describe them here, but full definitions are in appendix Sec. A.4. The sum functor \oplus denotes a functor that sends a cpo to the disjoint union \uplus of its component applications (with usual injection-wise order), and its functorial action is injection-wise. The product functor \otimes denotes a functor that sends a cpo to the product \times of its component applications (with usual component-wise order), and its functorial action is component-wise. The unit functor I denotes a functor sending any cpo to $1 = \{\bullet\}$ (discrete order), and its functorial action sends all morphisms to $id_{\{\bullet\}}$. The constant (type) functor \underline{P} denotes a functor sending any cpo to the cpo $\llbracket P \rrbracket$, and its functorial action sends all morphisms to the

identity $id_{\llbracket P \rrbracket}$ on $\llbracket P \rrbracket$. The identity functor Id denotes the identity endofunctor on \mathbf{Cpo} . (Forgetting the order structure, functors also denote endofunctors on the category \mathbf{Set} of sets and functions.)

We now explain the denotational semantics of our inductive types. Semantically, we build an inductive type (such as $\llbracket \text{List } A \rrbracket$), by repeatedly applying (the denotation of) its functor specification (such as $\llbracket \text{List } F_A \rrbracket$) to the initial object $\llbracket 0 \rrbracket = \emptyset$. For example,

$$\llbracket \text{List } A \rrbracket = \bigcup_{k \in \mathbb{N}} \llbracket \perp \oplus (A \otimes \text{Id}) \rrbracket^k \emptyset = 1 \uplus \left(\llbracket A \rrbracket \times \left(1 \uplus (\llbracket A \rrbracket \times \dots) \right) \right)$$

where $1 = \{\bullet\}$ (using the relatively direct functors with more complicated unrolling, discussed in Sec. 4.0.1). We denote the nil list $[]$ by $\text{inj}_1 \bullet$, a list $x :: []$ with one term x by $\text{inj}_2(\llbracket x \rrbracket, \text{inj}_1 \bullet)$, and so on. In general, given a (polynomial) \mathbf{Set} (category of sets and functions) endofunctor F (which, for this paper, will always be the denotation of a well-formed (syntactic) functor, refined or otherwise), we define $\mu F = \cup_{k \in \mathbb{N}} F^k \emptyset$. We then define $\llbracket \mu F \rrbracket = \mu \llbracket F \rrbracket$. In our system, for every well-formed (unrefined) functor F , the set $\mu \llbracket F \rrbracket$ is a *fixed point* of $\llbracket F \rrbracket$ (appendix Lemma C.7): that is, $\llbracket F \rrbracket (\mu \llbracket F \rrbracket) = \mu \llbracket F \rrbracket$ (and similarly for refined functors: appendix Lemma D.11).

Negative types. The grammar for unrefined negative types has unrefined function types $P \rightarrow N$ and unrefined upshifts $\uparrow P$, with no guarded or universal types. Unrefined negative types denote \mathbf{Cppo} s.

$$N ::= P \rightarrow N \mid \uparrow P$$

Function types $P \rightarrow N$ denote continuous functions from $\llbracket P \rrbracket$ to $\llbracket N \rrbracket$ (which we sometimes write as $\llbracket P \rrbracket \Rightarrow \llbracket N \rrbracket$), where its order is defined pointwise, together with the bottom element (the “point” of “pointed cpo”) $\perp_{\llbracket P \rightarrow N \rrbracket}$ that maps every $V \in \llbracket P \rrbracket$ to the bottom element $\perp_{\llbracket N \rrbracket}$ of $\llbracket N \rrbracket$ (that is, \uparrow denotes the lift functor from \mathbf{Cpo} to \mathbf{Cppo}). For our purposes, this is equivalent to lifting $\llbracket P \rrbracket \in \mathbf{Cpo}$ to \mathbf{Cppo} and denoting arrow types by *strict* (\perp goes to \perp) continuous functions so that function types denote \mathbf{Cppo} exponentials.

Upshifts $\uparrow P$ denote $\llbracket P \rrbracket \uplus \{\perp_{\uparrow}\}$ with the lift order

$$\sqsubseteq_{\llbracket \uparrow P \rrbracket} = \left\{ (\text{inj}_1 d, \text{inj}_1 d') \mid d \sqsubseteq_{\llbracket P \rrbracket} d' \right\} \cup \left\{ (\text{inj}_2 \perp_{\uparrow}, d) \mid d \in \llbracket \uparrow P \rrbracket \right\}$$

and bottom element $\perp_{\llbracket \uparrow P \rrbracket} = \text{inj}_2 \perp_{\uparrow}$. We could put, say, \bullet rather than \perp_{\uparrow} , but we think the latter is clearer in associating it with the *bottom* element of upshifts; or \perp rather than \perp_{\uparrow} but we often elide the “[A]” subscript in $\perp_{\llbracket A \rrbracket}$ when clear from context.

Appendix Fig. 27 has the full definition of (unrefined) type and functor denotations.

Well-typed program terms. We write $\Gamma \vdash O \dots A$ and $\Gamma; [B] \vdash O \dots A$ to stand for all six unrefined program typing judgments: $\Gamma \vdash h \Rightarrow P$ and $\Gamma \vdash g \Rightarrow \uparrow P$ and $\Gamma \vdash v \Leftarrow P$ and $\Gamma \vdash e \Leftarrow N$ and $\Gamma; [P] \vdash \{r_i \Rightarrow e_i\}_{i \in I} \Leftarrow N$ and $\Gamma; [N] \vdash s \gg \uparrow P$.

The denotational semantics of well-typed, unrefined program terms of judgmental form $\Gamma \vdash O \dots A$ or $\Gamma; [B] \vdash O \dots A$ are continuous functions $\llbracket \Gamma \rrbracket \rightarrow \llbracket A \rrbracket$ and $\llbracket \Gamma \rrbracket \rightarrow \llbracket B \rrbracket \rightarrow \llbracket A \rrbracket$ respectively, where $\llbracket \Gamma \rrbracket$ is the set of all semantic substitutions $\vdash \delta : \Gamma$ together with component-wise order. Similarly to function type denotations, the bottom element of a $\llbracket \Gamma \rrbracket \rightarrow \llbracket N \rrbracket$ sends every $\delta \in \llbracket \Gamma \rrbracket$ to $\perp_{\llbracket N \rrbracket}$ (equivalently for our purposes, we can lift source predomains and consider strict continuous functions). We only interpret typing derivations, but we often only mention the program term in semantic brackets $\llbracket - \rrbracket$. For example, if $\Gamma \vdash x \Rightarrow P$, then $\llbracket x \rrbracket = (\delta \in \llbracket \Gamma \rrbracket) \mapsto \delta(x)$. We write the application of the denotation $\llbracket E \rrbracket$ of a program term E (typed under Γ) to a semantic substitution $\delta \in \llbracket \Gamma \rrbracket$ as $\llbracket E \rrbracket_{\delta}$. We only mention a few of the more interesting cases of the definition of typing denotations; for the full definition, see appendix Figures 28, 29, and 30. If $\Gamma; [N] \vdash v, s \gg M$, then

$$\llbracket v, s \rrbracket = (\delta \in \llbracket \Gamma \rrbracket) \mapsto (f \mapsto \llbracket s \rrbracket_{\delta} (f(\llbracket v \rrbracket_{\delta})))$$

Returner expressions denote monadic returns:

$$\llbracket \text{return } v \rrbracket_{\delta} = \text{inj}_1 \llbracket v \rrbracket_{\delta}$$

Let-binding denotes monadic binding:

$$\llbracket \text{let } x = g; e \rrbracket_{\delta} = \begin{cases} \llbracket e \rrbracket_{(\delta, V/x)} & \text{if } \llbracket g \rrbracket_{\delta} = \text{inj}_1 V \\ \perp_{\llbracket N \rrbracket} & \text{if } \llbracket g \rrbracket_{\delta} = \text{inj}_2 \perp_{\uparrow} \end{cases}$$

A recursive expression denotes a fixed point obtained by taking the least upper bound (\sqcup) of all its successive approximations:

$$\llbracket \Gamma \vdash \text{rec } x. e \Leftarrow N \rrbracket_{\delta} = \bigsqcup_{k \in \mathbb{N}} \left(V \mapsto \llbracket \Gamma, x : \downarrow N \vdash e \Leftarrow N \rrbracket_{\delta, V/x} \right)^k \perp_{\llbracket N \rrbracket}$$

In the unrefined system, we include `diverge`, to which `unreachable` erases (that is, $|\text{unreachable}| = \text{diverge}$). We intend `diverge` to stand for an undefined body of a pattern-matching clause, but we interpret this error as divergence to simplify the semantics:

$$\llbracket \Gamma \vdash \text{diverge} \Leftarrow N \rrbracket_{\delta} = \perp_{\llbracket N \rrbracket}$$

The point is that the refined system prevents the error.

We will say more about the semantics of folds in Sec. 5.2, but note that the action of rolling and unrolling syntactic values is essentially denoted by $d \mapsto d$:

$$\begin{aligned} \llbracket \text{into}(v) \rrbracket_{\delta} &= \llbracket v \rrbracket_{\delta} \\ \llbracket \{\text{into}(x) \Rightarrow e\} \rrbracket_{\delta} &= V \mapsto \llbracket e \rrbracket_{\delta, V/x} \end{aligned}$$

This works due to the fact that unrolling is sound (roughly, the denotations of each side of “ \doteq ” in the unrolling judgment are *equal*) and the fact that $\llbracket F \rrbracket (\mu \llbracket F \rrbracket) = \mu \llbracket F \rrbracket$ (and similarly for the refined system).

Unrefined soundness. Our proofs of (appendix) Lemma C.28 (Unrefined Type Soundness) and (appendix) Lemma C.30 (Unrefined Substitution Soundness) use standard techniques in domain theory [Gunter 1993].

Unrefined type soundness says that a term typed A under Γ denotes a continuous function $\llbracket \Gamma \rrbracket \rightarrow \llbracket A \rrbracket$. We (partly) state (3 out of 6 parts) this in two mutually recursive lemmas as follows:

LEMMA 5.1 (CONTINUOUS MAPS).

(Lemma C.27 in appendix)

Suppose $\vdash \delta_1 : \Gamma_1$ and $\vdash \delta_2 : \Gamma_2$ and $\vdash \Gamma_1, y : Q, \Gamma_2$ ctx.

- (1) If $\Gamma_1, y : Q, \Gamma_2 \vdash h \Rightarrow P$, then the function $\llbracket Q \rrbracket \rightarrow \llbracket P \rrbracket$ defined by $d \mapsto \llbracket h \rrbracket_{\delta_1, d/y, \delta_2}$ is continuous.
- (2) If $\Gamma_1, y : Q, \Gamma_2 \vdash e \Leftarrow N$, then the function $\llbracket Q \rrbracket \rightarrow \llbracket N \rrbracket$ defined by $d \mapsto \llbracket e \rrbracket_{\delta_1, d/y, \delta_2}$ is continuous.
- (3) If $\Gamma_1, y : Q, \Gamma_2; \llbracket N \rrbracket \vdash s \gg \uparrow P$, then the function $\llbracket Q \rrbracket \rightarrow \llbracket N \rrbracket \rightarrow \llbracket \uparrow P \rrbracket$ defined by $d \mapsto \llbracket s \rrbracket_{\delta_1, d/y, \delta_2}$ is continuous.

LEMMA 5.2 (UNREFINED TYPE SOUNDNESS).

(Lemma C.28 in appendix)

Assume $\vdash \delta : \Gamma$.

- (1) If $\Gamma \vdash h \Rightarrow P$, then $\llbracket \Gamma \vdash h \Rightarrow P \rrbracket_{\delta} \in \llbracket P \rrbracket$.
- (2) If $\Gamma \vdash e \Leftarrow N$, then $\llbracket \Gamma \vdash e \Leftarrow N \rrbracket_{\delta} \in \llbracket N \rrbracket$.
- (3) If $\Gamma; \llbracket N \rrbracket \vdash s \gg \uparrow P$, then $\llbracket \Gamma; \llbracket N \rrbracket \vdash s \gg \uparrow P \rrbracket_{\delta} \in \llbracket N \rrbracket \Rightarrow \llbracket \uparrow P \rrbracket$.

The proof of unrefined type soundness is standard, and uses the well-known fact that a continuous function in **Cppo** has a least fixed point. Among other things, we also use the fact that $\mu \llbracket F \rrbracket$ is a fixed point of $\llbracket F \rrbracket$ (appendix Lemma C.7). We also use the soundness of unrefined unrolling, which

we didn't mention here because it's similar to refined unrolling and its soundness, discussed in the next section.

We interpret an unrefined syntactic substitution (typing derivation) $\Gamma_0 \vdash \sigma : \Gamma$ as a continuous function $\llbracket \Gamma_0 \rrbracket \rightarrow \llbracket \Gamma \rrbracket$ that takes a $\delta \in \llbracket \Gamma_0 \rrbracket$ and uses δ to interpret each of the entries in σ (remembering to apply the rest of the syntactic substitution, because substitution is defined sequentially):

$$\begin{aligned} \llbracket \Gamma_0 \vdash \cdot : \cdot \rrbracket &= (\delta \in \llbracket \Gamma_0 \rrbracket) \mapsto \cdot \\ \llbracket \Gamma_0 \vdash (\sigma, (v : P/x)) : (\Gamma, x : P) \rrbracket &= (\delta \in \llbracket \Gamma_0 \rrbracket) \mapsto ((\llbracket \sigma \rrbracket)_\delta, \llbracket [\sigma]v \rrbracket_\delta / x) \end{aligned}$$

Similarly to typing derivations, we only consider denotations of typing derivations $\Gamma_0 \vdash \sigma : \Gamma$ of substitutions, but often simply write $\llbracket \sigma \rrbracket$.

Unrefined substitution soundness says that semantic and syntactic substitution commute: if E is a program term typed under Γ and $\Gamma_0 \vdash \sigma : \Gamma$ is a substitution, then $\llbracket [\sigma]E \rrbracket = \llbracket E \rrbracket \circ \llbracket \sigma \rrbracket$. Here, we partly show how it is stated in the appendix (1 out of 6 parts):

LEMMA 5.3 (UNREFINED SUBSTITUTION SOUNDNESS).

(Lemma C.30 in appendix)

Assume $\Gamma_0 \vdash \sigma : \Gamma$ and $\vdash \delta : \Gamma_0$.

(1) If $\Gamma \vdash e \leftarrow N$, then $\llbracket \Gamma_0 \vdash [\sigma]e \leftarrow N \rrbracket_\delta = \llbracket \Gamma \vdash e \leftarrow N \rrbracket_{\llbracket \sigma \rrbracket_\delta}$.

We use unrefined type/substitution soundness to prove *refined* type/substitution soundness, discussed next.

5.2 Refined System

Indexes. For any sort τ , we give its denotation $\llbracket \tau \rrbracket$ the discrete order $\sqsubseteq_{\llbracket \tau \rrbracket} = \{(d, d) \mid d \in \llbracket \tau \rrbracket\}$, making it a cpo.

Semantic Substitution. We introduced semantic substitutions δ (at the index level) when discussing propositional validity (Sec. 4.1). Here, they are extended to semantic program values:

$$\frac{\vdash \delta : \Theta; \Gamma \quad V \in \llbracket P \rrbracket_{\llbracket \delta \rrbracket} \quad x \notin \text{dom}(\Gamma)}{\vdash (\delta, V/x) : \Theta; \Gamma, x : P}$$

where $\llbracket - \rrbracket$ filters out program entries. *Notation:* we define $\llbracket \Theta; \Gamma \rrbracket = \{\delta \mid \vdash \delta : \Theta; \Gamma\}$.

Erasure. The *erasure* metaoperation $|-|$ (appendix Sec. A.5) erases all indexes from (refined) types, program terms (which can have type annotations, but those do not affect program meaning), and syntactic and semantic substitutions. For example, $|\{v : \mu F \mid (\text{fold}_F \alpha) v =_\tau t\}| = \mu|F|$ and $|\forall a : \tau. N| = |N|$ and $|P \times Q| = |P| \times |Q|$ and so on.

We use many facts about erasure to prove refined type/substitution soundness (appendix lemmas):

- Refined types denote subsets of what their erasures denote: Lemma C.31 (Type Subset of Erasure). Similarly for refined functors and refined inductive types: Lemma C.32 (Functor Application Subset of Erasure) and Lemma C.33 (Mu Subset of Erasure).
- The erasure of both types appearing in extraction, equivalence, and subtyping judgments results in equal (unrefined) types: Lemma C.36 (Extraction Erases to Equality), Lemma C.37 (Equivalence Erases to Equality), and Lemma C.38 (Subtyping Erases to Equality).
- Refined unrolling and typing are sound with respect to their erasure: Lemma C.39 (Erasure Respects Unrolling), Lemma C.40 (Erasure Respects Typing), and Lemma C.42 (Erasure Respects Substitution Typing).
- Erasure commutes with syntactic and semantic substitution: Lemma C.41 (Erasure Respects Substitution) and Lemma C.43 (Erasure Respects Semantic Substitution).

Types, functors, algebras, and folds. The denotations of refined types and functors are defined as logical subsets of the denotations of their erasures (together with their erasure denotations themselves). They are defined mutually with the denotations of well-formed algebras.

In appendix Fig. 36, we inductively define the denotations of well-formed types $\Theta \vdash A$ type $[_]$. We briefly discuss a few of the cases. The meaning of an asserting type is the set of refined values such that the asserted index proposition holds (read $\{\bullet\}$ as true and \emptyset as false):

$$\llbracket P \wedge \phi \rrbracket_\delta = \{V \in \llbracket |P| \rrbracket \mid V \in \llbracket P \rrbracket_\delta \text{ and } \llbracket \phi \rrbracket_\delta = \{\bullet\}\}$$

Existential and universal types denote elements of their erasure such that the relevant index quantification holds:

$$\begin{aligned} \llbracket \exists a : \tau. P \rrbracket_\delta &= \left\{ V \in \llbracket |P| \rrbracket \mid \exists d \in \llbracket \tau \rrbracket. V \in \llbracket P \rrbracket_{\delta, d/a} \right\} \\ \llbracket \forall a : \tau. N \rrbracket_\delta &= \left\{ f \in \llbracket |N| \rrbracket \mid \forall d \in \llbracket \tau \rrbracket. f \in \llbracket N \rrbracket_{\delta, d/a} \right\} \end{aligned}$$

Guarded types denote elements of their erasure such that they are also in the refined type being guarded if the guard holds ($\{\bullet\}$ means true):

$$\llbracket \phi \supset N \rrbracket_\delta = \{f \in \llbracket |N| \rrbracket \mid \text{if } \llbracket \phi \rrbracket_\delta = \{\bullet\} \text{ then } f \in \llbracket N \rrbracket_\delta\}$$

The denotation of refined function types $\llbracket P \rightarrow N \rrbracket_\delta$ is *not* the set $\llbracket P \rrbracket_\delta \Rightarrow \llbracket N \rrbracket_\delta$ of (continuous) functions from refined P -values to refined N -values; if it were, then type soundness would break:

$$\llbracket \cdot; \cdot \vdash \lambda x. \text{return } x \Leftarrow (1 \wedge \text{ff}) \rightarrow \uparrow 1 \rrbracket. = (\bullet \mapsto \text{inj}_1 \bullet)$$

which is not in $(\emptyset \Rightarrow \{\bullet\} \uplus \{\perp_\uparrow\})$. Instead, the meaning of a refined function type is a set (resembling a unary logical relation)

$$\{f \in \llbracket |P \rightarrow N| \rrbracket \mid \forall V \in \llbracket P \rrbracket_\delta. f(V) \in \llbracket N \rrbracket_\delta\}$$

of *unrefined* (continuous) functions that take refined values to refined values. The meaning of *refined* upshifts enforces termination (if refined type soundness holds, and we will see it does):

$$\llbracket \uparrow P \rrbracket_\delta = \{\text{inj}_1 V \mid V \in \llbracket P \rrbracket_\delta\}$$

Note that divergence $\text{inj}_2 \perp_\uparrow$ is *not* in the set $\llbracket \uparrow P \rrbracket_\delta$.

In appendix Fig. 37, we inductively define the denotations of well-formed refined functors F and algebras α . The main difference between refined and unrefined functors is that in refined functors, constant functors produce subsets of their erasure. All functors, refined or otherwise, also (forgetting the partial order structure) denote endofunctors on the category of sets and functions. As with our unrefined functors, our refined functors denote functors with a fixed point (appendix Lemma D.11): $\llbracket F \rrbracket_\delta (\mu \llbracket F \rrbracket_\delta) = \mu \llbracket F \rrbracket_\delta$. Moreover, $\mu \llbracket F \rrbracket_\delta$ satisfies a recursion principle such that we can inductively define measures on $\mu \llbracket F \rrbracket_\delta$ via $\llbracket F \rrbracket_\delta$ -algebras (discussed next).

Categorically, given an endofunctor F , an F -algebra is an evaluator map $\alpha : F(\tau) \rightarrow \tau$ for some carrier set τ . We may think of this in terms of elementary algebra: we form algebraic expressions with F and evaluate them with α . A morphism f from algebra $\alpha : F(\tau) \rightarrow \tau$ to algebra $\beta : F(\tau') \rightarrow \tau'$ is a morphism $f : \tau \rightarrow \tau'$ such that $f \circ \alpha = \beta \circ (F(f))$. If an endofunctor F has an *initial*⁷ algebra *into* $: F(\mu F) \rightarrow \mu F$, then it has a recursion principle. By the recursion principle for μF , we can define a recursive function from μF to τ by *folding* μF with an F -algebra $\alpha : F(\tau) \rightarrow \tau$ like so:

$$\begin{aligned} (\text{fold}_F \alpha) : \mu F &\rightarrow \tau \\ (\text{fold}_F \alpha) v &= \alpha \left((\text{fmap } F (\text{fold}_F \alpha)) (\text{out}(v)) \right) \end{aligned}$$

⁷An object X in a category \mathbf{C} is *initial* if for every object Y in \mathbf{C} , there exists a unique morphism $X \rightarrow Y$ in \mathbf{C} .

where $out : \mu F \rightarrow F(\mu F)$, which by Lambek’s lemma exists and is inverse to $into$, embeds (semantic) inductive values into the unrolling of the (semantic) inductive type (we usually elide $fmap$). Conveniently, in our system and semantics, out is always $d \mapsto d$, and we almost never explicitly mention it. Syntactic values v in our system must be rolled into inductive types— $into(v)$ —and this is also how (syntactic) inductive values are pattern-matched (“applying out ” to $into(v)$), but $into(-)$ conveniently denotes $d \mapsto d$.

We specify inductive types abstractly as sums of products so that they denote polynomial endofunctors more directly. Polynomial endofunctors always have a “least” (initial) fixed point⁸, and hence specify inductive types, which have a recursion principle. For example, we specify (modulo the unrolling simplification) $len : List_{F_A}(\mathbb{N}) \Rightarrow \mathbb{N}$ (Sec. 1) by the (syntactic) algebra

$$\alpha = inj_1 () \Rightarrow 0 \mid inj_2 (\top, a) \Rightarrow 1 + a$$

which denotes the (semantic) algebra

$$\llbracket \alpha \rrbracket : \underbrace{\llbracket List_{F_A} \rrbracket (\mathbb{N})}_{1 \uplus (\llbracket A \rrbracket \times \mathbb{N})} \rightarrow \mathbb{N}$$

defined by $\llbracket \alpha \rrbracket = [\bullet \mapsto 0, (a, n) \mapsto 1 + n]$. By initiality (the recursion principle), there is a unique function

$$fold_{\llbracket List_{F_A} \rrbracket} \llbracket \alpha \rrbracket : \mu \llbracket List_{F_A} \rrbracket \rightarrow \mathbb{N}$$

such that $fold_{\llbracket List_{F_A} \rrbracket} \llbracket \alpha \rrbracket = \llbracket \alpha \rrbracket \circ ((F) (fold_{\llbracket List_{F_A} \rrbracket} \llbracket \alpha \rrbracket))$, which semantically captures len (Sec. 1).

In our system, a refined inductive type is written $\{v : \mu F \mid (fold_F \alpha) v =_{\tau} t\}$, which looks quite similar to its own semantics:

$$\llbracket \{v : \mu F \mid (fold_F \alpha) v =_{\tau} t\} \rrbracket_{\delta} = \left\{ V \in \mu \llbracket F \rrbracket_{\delta} \mid (fold_{\llbracket F \rrbracket_{\delta}} \llbracket \alpha \rrbracket_{\delta}) V = \llbracket t \rrbracket_{\delta} \right\}$$

The type $List(A)(n)$ of A -lists having length $n : \mathbb{N}$, for example, is defined in our system as:

$$List(A)(n) = \{v : \mu List_{F_A} \mid (fold_{List_{F_A}} len_{\text{alg}}) v =_{\mathbb{N}} n\}$$

Syntactic types, functors, and algebras in our system look very similar to their own semantics.

A well-typed algebra $\Xi; \Theta \vdash \alpha : F(\tau) \Rightarrow \tau$ denotes a dependent function $\prod_{\delta \in [\Theta]} \llbracket F \rrbracket_{\delta} \llbracket \tau \rrbracket \rightarrow \llbracket \tau \rrbracket$. The definition (appendix Fig. 37) is mostly standard, but the unit and pack cases could use some explanation. Because Θ is for F and $\Xi (\subseteq \Theta)$ is for α , we restrict δ to Ξ at algebra bodies:

$$\llbracket \Xi; \Theta \vdash () \Rightarrow t : I(\tau) \Rightarrow \tau \rrbracket_{\delta} \bullet = \llbracket \Xi \vdash t : \tau \rrbracket_{\delta \upharpoonright \Xi}$$

The most interesting part of the definition concerns index packing:

$$\begin{aligned} \llbracket \Xi; \Theta \vdash (\text{pack}(a, o), q) \Rightarrow t : (\exists a : \tau'. Q \otimes \hat{P})(\tau) \Rightarrow \tau \rrbracket_{\delta} (V_1, V_2) = \\ \llbracket \Xi, a : \tau'; \Theta, a : \tau' \vdash (o, q) \Rightarrow t : (Q \otimes \hat{P})(\tau) \Rightarrow \tau \rrbracket_{(\delta, d/a)} (V_1, V_2) \\ \text{where } d \in \llbracket \tau' \rrbracket \text{ is such that } V_1 \in \llbracket Q \rrbracket_{\delta, d/a} \end{aligned}$$

The pack clause lets us bind the witness d of τ' in the existential type $\exists a : \tau'. Q$ to a in the body t of the algebra. We know d exists since $V_1 \in \llbracket \exists a : \tau'. Q \rrbracket_{\delta}$, but it is not immediate that it is unique. However, we prove d is uniquely determined by V_1 ; we call this property the *soundness of value-determined indexes* (all parts are mutually recursive):

LEMMA 5.4 (SOUNDNESS OF VALUE-DETERMINED INDEXES).

(Lemma D.14 in appendix)

Assume $\vdash \delta_1 : \Theta$ and $\delta_2 : \Theta$.

⁸This is not the case for all endofunctors. Therefore, not all endofunctors can be said to specify an inductive type. For example, consider the powerset functor.

- (1) If $\Theta \vdash P \text{ type}[\Xi]$ and $V \in \llbracket P \rrbracket_{\delta_1}$ and $V \in \llbracket P \rrbracket_{\delta_2}$, then $\delta_1 \upharpoonright_{\Xi} = \delta_2 \upharpoonright_{\Xi}$.
- (2) If $\Theta \vdash \mathcal{F} \text{ functor}[\Xi]$ and $X_1, X_2 \in \text{Set}$ and $V \in \llbracket \mathcal{F} \rrbracket_{\delta_1} X_1$ and $V \in \llbracket \mathcal{F} \rrbracket_{\delta_2} X_2$, then $\delta_1 \upharpoonright_{\Xi} = \delta_2 \upharpoonright_{\Xi}$.
- (3) If $\Xi; \Theta \vdash \alpha : F(\tau) \Rightarrow \tau$ and $\Xi \subseteq \Theta$ and $\delta_1 \upharpoonright_{\Xi} = \delta_2 \upharpoonright_{\Xi}$, then $\llbracket \alpha \rrbracket_{\delta_1} = \llbracket \alpha \rrbracket_{\delta_2}$ on $\llbracket F \rrbracket_{\delta_1} \llbracket \tau \rrbracket \cap \llbracket F \rrbracket_{\delta_2} \llbracket \tau \rrbracket$.

Therefore, the Ξ in type and functor well-formedness really does track index variables that are uniquely determined by values, semantically speaking.

Well-typed program terms. Appendix Fig. 38 specifies the denotations of well-typed refined program terms in terms of the denotations of their erasure. The denotation of a refined program term E typed under $(\Theta; \Gamma)$, at refined semantic substitution $\delta \in \llbracket \Theta; \Gamma \rrbracket$, is the denotation $\llbracket |E| \rrbracket_{|\delta|}$ of the (derivation of the) term's erasure $|E|$ at the erased substitution $|\delta|$. For example,

$$\llbracket \Theta; \Gamma \vdash e \leftarrow N \rrbracket = (\delta \in \llbracket \Theta; \Gamma \rrbracket) \mapsto \llbracket |\Gamma| \vdash |e| \leftarrow |N| \rrbracket_{|\delta|}$$

Unrolling. We prove (appendix Lemma D.15) that unrolling is sound:

LEMMA 5.5 (UNROLLING SOUNDNESS). (Lemma D.15 in appendix)

Assume $\vdash \delta : \Theta$ and $\Xi \subseteq \Theta$. If $\Xi; \Theta \vdash \{v : G[\mu F] \mid \beta(G(\text{fold}_F \alpha) v) =_\tau t\} \doteq P$, then $\left\{ V \in \llbracket G \rrbracket_{\delta} (\mu \llbracket F \rrbracket_{\delta}) \mid \llbracket \beta \rrbracket_{\delta} (\llbracket G \rrbracket_{\delta} (\text{fold}_{\llbracket F \rrbracket_{\delta}} \llbracket \alpha \rrbracket_{\delta}) V) = \llbracket t \rrbracket_{\delta} \right\} = \llbracket P \rrbracket_{\delta}$.

Due to our definition of algebra denotations (specifically, for the pack pattern), we use the soundness of value-determined indexes in the pack case of the proof.

Subtyping. We prove (appendix Lemma D.19) that subtyping is sound:

LEMMA 5.6 (SOUNDNESS OF SUBTYPING). (Lemma D.19 in appendix)

Assume $\vdash \delta : \Theta$. If $\Theta \vdash A \leq^{\pm} B$, then $\llbracket A \rrbracket_{\delta} \subseteq \llbracket B \rrbracket_{\delta}$.

Type soundness. Denotational-semantic type soundness says that if a program term has type A under Θ and Γ , then the mathematical meaning of that program term at any interpretation of (that is, semantic environment for) Θ and Γ is an element of the mathematical meaning of A at that interpretation, that is, the program term denotes a dependent function $\prod_{\delta \in \llbracket \Theta; \Gamma \rrbracket} \llbracket A \rrbracket_{|\delta|}$. This more or less corresponds to proving (operational) type soundness with respect to a big-step operational semantics. Refined types pick out subsets of values of unrefined types. Therefore, by type soundness, if a program has a refined type, then we have learned something more about that program than the unrefined system can verify for us.

THEOREM 5.7 (TYPE SOUNDNESS). (Thm. D.25 in appendix)

Assume $\vdash \delta : \Theta; \Gamma$. Then:

- (1) If $\Theta; \Gamma \vdash h \Rightarrow P$, then $\llbracket h \rrbracket_{\delta} \in \llbracket P \rrbracket_{|\delta|}$.
- (2) If $\Theta; \Gamma \vdash g \Rightarrow N$, then $\llbracket g \rrbracket_{\delta} \in \llbracket N \rrbracket_{|\delta|}$.
- (3) If $\Theta; \Gamma \vdash v \leftarrow P$, then $\llbracket v \rrbracket_{\delta} \in \llbracket P \rrbracket_{|\delta|}$.
- (4) If $\Theta; \Gamma \vdash e \leftarrow N$, then $\llbracket e \rrbracket_{\delta} \in \llbracket N \rrbracket_{|\delta|}$.
- (5) If $\Theta; \Gamma; [P] \vdash \{r_i \Rightarrow e_i\}_{i \in I} \leftarrow N$, then $\llbracket \{r_i \Rightarrow e_i\}_{i \in I} \rrbracket_{\delta} \in \llbracket P \rrbracket_{|\delta|} \Rightarrow \llbracket N \rrbracket_{|\delta|}$.
- (6) If $\Theta; \Gamma; [N] \vdash s \gg \uparrow P$, then $\llbracket s \rrbracket_{\delta} \in \llbracket N \rrbracket_{|\delta|} \Rightarrow \llbracket \uparrow P \rrbracket_{|\delta|}$.

(All parts are mutually recursive.) The proof (appendix Thm. D.25) uses the soundness of unrolling and subtyping. The proof is mostly straightforward. The hardest case is the one for recursive expressions in part (4), where we use an *upward closure* lemma—in particular, part (3) below—to show that the fixed point is in the appropriately refined set:

LEMMA 5.8 (UPWARD CLOSURE). (Lemma D.22 in appendix)

Assume $\vdash \delta : \Theta$.

- (1) If $\Xi; \Theta \vdash \alpha : F(\tau) \Rightarrow \tau$ and $\Xi \subseteq \Theta$ then $\llbracket \alpha \rrbracket_\delta$ is monotone.
- (2) If $\Theta \vdash \mathcal{G}$ functor $\llbracket _ \rrbracket$ and $\Theta \vdash F$ functor $\llbracket _ \rrbracket$ and $k \in \mathbb{N}$
and $V \in \llbracket \mathcal{G} \rrbracket_\delta (\llbracket F \rrbracket_\delta^k \emptyset)$ and $V \sqsubseteq_{\llbracket \mathcal{G} \rrbracket (\llbracket F \rrbracket^k \emptyset)} V'$,
then $V' \in \llbracket \mathcal{G} \rrbracket_\delta (\llbracket F \rrbracket_\delta^k \emptyset)$.
- (3) If $\Theta \vdash A$ type $\llbracket _ \rrbracket$ and $V \in \llbracket A \rrbracket_\delta$ and $V \sqsubseteq_{\llbracket A \rrbracket} V'$, then $V' \in \llbracket A \rrbracket_\delta$.

Out of all proofs in this paper, the proof of upward closure (appendix Lemma D.22) is a top contender for the most interesting induction metric:

PROOF. By lexicographic induction on, first, $\text{sz}(A)/\text{sz}(F)$ (parts (1), (2) and (3), mutually), and, second, $\langle k, \mathcal{G} \text{ structure} \rangle$ (part (2)), where $\langle \dots \rangle$ denotes lexicographic order. \square

We define the simple size function $\text{sz}(-)$, which is basically a standard structural notion of size, in appendix Fig. 57. This is also the only place, other than unrolling soundness, where we use the soundness of value-determined indexes (again for a pack case, in part (1)).

Substitution soundness. We interpret a syntactic substitution (typing derivation) $\Theta_0; \Gamma_0 \vdash \sigma : \Theta; \Gamma$ as a function $\llbracket \sigma \rrbracket : \llbracket \Theta_0; \Gamma_0 \rrbracket \rightarrow \llbracket \Theta; \Gamma \rrbracket$ on semantic substitutions (appendix Def. B.1). Similarly to the interpretation of unrefined substitution typing derivations, the interpretation of the head term being substituted (its typing/sorting subderivation) pre-applies the rest of the substitution:

$$\begin{aligned} \llbracket \cdot \rrbracket_\delta &= \cdot \\ \llbracket \sigma, t/a \rrbracket_\delta &= \llbracket \sigma \rrbracket_\delta, \llbracket \llbracket \sigma \rrbracket t \rrbracket_{\llbracket \delta \rrbracket} / a \\ \llbracket \sigma, v : P/x \rrbracket_\delta &= \llbracket \sigma \rrbracket_\delta, \llbracket \llbracket \sigma \rrbracket v \rrbracket_\delta / x \end{aligned}$$

Substitution soundness holds (appendix Thm. D.28): if E is a program term typed under Θ and Γ , and $\Theta_0; \Gamma_0 \vdash \sigma : \Theta; \Gamma$, then $\llbracket \llbracket \sigma \rrbracket E \rrbracket = \llbracket E \rrbracket \circ \llbracket \sigma \rrbracket$. (Recall we prove a syntactic substitution lemma: Lemma 4.2.) That is, substitution and denotation commute, or (in other words) syntactic substitution and semantic substitution are compatible.

Logical consistency, total correctness, and partial correctness. Our *semantic* type soundness result implies that our system is logically consistent and totally correct.

A logically inconsistent type (for example, 0 or $\uparrow 0$ or $\uparrow (1 \wedge \text{ff})$) denotes the empty set, which is uninhabited.

COROLLARY 5.9 (LOGICAL CONSISTENCY). *If $\cdot; \cdot \vdash e \Leftarrow N$, then N is logically consistent, that is, $\llbracket N \rrbracket \neq \emptyset$. Similarly, if $\cdot; \cdot \vdash v \Leftarrow P$, then P is logically consistent, and so on for the other typing judgments.*

Proving logical consistency *syntactically*, say, via progress and preservation lemmas, would require *also* proving that every reduction sequence eventually terminates (that is, strong normalization), which might need a relatively complicated proof using logical relations [Tait 1967].

Total correctness means that every closed computation (that is specified as total) returns a value of the specified type:

COROLLARY 5.10 (TOTAL CORRECTNESS). *If $\cdot; \cdot \vdash e \Leftarrow \uparrow P$, then $\llbracket e \rrbracket \neq \perp_{\llbracket P \rrbracket}$, that is, e does not diverge.*

PROOF.

$$\begin{aligned} \vdash \cdot; \cdot & && \text{By Empty}\delta \\ \llbracket e \rrbracket \in \llbracket \llbracket \uparrow P \rrbracket \rrbracket_{\llbracket \cdot \rrbracket} & && \text{By Theorem 5.7 (Type Soundness)} \\ = \llbracket \llbracket \uparrow P \rrbracket \rrbracket & && \text{By definition of } \llbracket _ \rrbracket \\ = \{ \text{inj}_1 V \mid V \in \llbracket P \rrbracket \} & && \text{By definition of } \llbracket _ \rrbracket. \end{aligned}$$

Therefore, $\llbracket e \rrbracket. \neq inj_2 \perp_{\uparrow} = \perp_{\llbracket \uparrow P \rrbracket}$, that is, e terminates (and returns a value). \square

Our system can be extended to include partiality, simply by adding a *partial* upshift type connective $\uparrow P$ (“partial upshift of P ”), with type well-formedness, subtyping and type equivalence rules similar to those of $\uparrow P$, and the following two expression typechecking rules. The first rule introduces the new connective $\uparrow P$; the second rule lacks a termination refinement such as that in $\text{Decl} \Leftarrow \text{rec}$, so it may yield divergence.

$$\frac{\Theta; \Gamma \vdash v \Leftarrow P}{\Theta; \Gamma \vdash \text{return } v \Leftarrow \uparrow P} \qquad \frac{\Theta; \Gamma, x : \downarrow N \vdash e \Leftarrow N}{\Theta; \Gamma \vdash \text{rec } x. e \Leftarrow N}$$

The meaning of the partial upshift is defined as follows:

$$\llbracket \Theta \vdash \uparrow P \text{ type } [_]\rrbracket_{\delta} = \{d \in \llbracket \uparrow P \rrbracket \mid \text{if } d \neq \perp_{\llbracket \uparrow P \rrbracket} \text{ then } d = inj_1 V \text{ for some } V \in \llbracket P \rrbracket_{\delta}\}$$

It is straightforward to update the metatheory to prove *partial correctness*: If a closed computation (that is specified as partial) terminates, then it returns a value of the specified type. Partial correctness is a corollary of the updated type soundness result: if $\cdot \vdash e \Leftarrow \uparrow P$ and $\llbracket e \rrbracket. \neq \perp_{\llbracket \uparrow P \rrbracket}$, then $\llbracket e \rrbracket. = inj_1 V$ and $V \in \llbracket P \rrbracket.$

Adding partiality introduces logical inconsistency, so we must restate logical consistency for expression typing: If $\cdot \vdash e \Leftarrow \uparrow P$, then $\uparrow P$ is *logically consistent*.

6 ALGORITHMIC SYSTEM

We design our algorithmic system in the spirit of those of Dunfield and Krishnaswami [2013, 2019], but those systems do not delay constraint verification until all existentials are solved. The algorithmic rules closely mirror the declarative rules, except for a few key differences:

- Whenever a declarative rule conjures an index term, the corresponding algorithmic rule adds, to a separate (input) algorithmic context Δ , an existential variable (written with a hat: \hat{a}) to be solved.
- As the typing algorithm proceeds, we add index term solutions of the existential variables to the output algorithmic context, increasing knowledge (see Sec. 6.2). We eagerly apply index solutions to input types and output constraints, and pop them off the output context when out of scope.
- Whenever a declarative rule checks propositional validity or equivalence ($\Theta \vdash \phi$ true or $\Theta \vdash \phi \equiv \psi$), the algorithm delays checking the constraint until all existential variables in the propositions are solved (at the end of a focusing stage). Similarly, subtyping, type equivalence, and expression typechecking constraints are delayed until all existential variables are solved. When an entity has no existential variables, we say that it is *ground*.
- In subtyping, we eagerly extract from assumptive positions immediately under polarity shifts.

Syntactically, objects in the algorithmic system are not much different from corresponding objects of the declarative system. We extend the grammar for index terms with a production of existential variables, which we write as an index variable with a hat \hat{a} , \hat{b} , or \hat{c} :

$$t ::= \dots \mid \hat{a}$$

We use this (algorithmic) index grammar everywhere in the algorithmic system, using the same declarative metavariables. However, we write algorithmic logical contexts with a hat: $\hat{\Theta}$. Algorithmic logical contexts $\hat{\Theta}$ only appear in output mode, and are like (input) logical contexts Θ , but propositions listed in them may have existential variables (its index variable sortings $a : \tau$ are universal).

Constraints are added to the algorithmic system. Figure 14 gives grammars for subtyping and typing constraints. In contrast to DML, the grammar does not include existential constraints.

$$\begin{array}{l}
 \text{Subtyping constraints } W ::= \phi \mid \phi \equiv \psi \mid \phi \supset W \mid W \wedge W \mid \forall a : \tau. W \\
 \quad \mid \underline{P < :^+ Q} \mid \underline{N < :^- M} \mid P \equiv^+ Q \mid N \equiv^- M \\
 \text{Typing constraints } \chi ::= \cdot \mid (e \leftarrow N), \chi \mid W, \chi
 \end{array}$$

Fig. 14. Typing and subtyping constraints

Checking constraints boils down to checking propositional validity, $\Theta \vdash \phi$ true, which is analogous to checking *verification conditions* in the tradition of imperative program verification initiated by Floyd [1967] and Hoare [1969] (where programs annotated with Floyd–Hoare assertions are analyzed, generating verification conditions whose validity implies program correctness). These propositional validity constraints are the constraints that can be automatically verified by a theorem prover such as an SMT solver. The (algorithmic) W constraint verification judgment is written $\Theta \models W$ and means that W algorithmically holds under Θ . Notice that the only context in the judgment is Θ , which has no existential variables: this reflects the fact that we delay verifying W until W has no existential variables (in which case we say W is *ground*). Similarly, $\Theta; \Gamma \triangleleft \chi$ is the (algorithmic) χ verification judgment, meaning all of the constraints in χ algorithmically hold under Θ and Γ , and here χ is also ground (by focusing).

6.1 Contexts and Substitution

Algorithmic contexts Δ are lists of solved or unsolved existential variables, and are said to be *complete*, and are written as Ω , if they are all solved:

$$\begin{array}{l}
 \Delta ::= \cdot \mid \Delta, \hat{a} : \tau \mid \Delta, \hat{a} : \tau = t \\
 \Omega ::= \cdot \mid \Omega, \hat{a} : \tau = t
 \end{array}$$

We require solutions t of existential variables \hat{a} to be well-sorted under (input) logical contexts Θ , which have no existential variables. To maintain this invariant that every solution in Δ is *ground*, that is, has no existential variables, we exploit type polarity in algorithmic subtyping, and prevent existential variables from ever appearing in refinement algebras.

We will often treat algorithmic contexts Δ as substitutions of ground index terms for existential variables \hat{a} in index terms t (including propositions ϕ), types A , functors \mathcal{F} , constraints W and χ , and output logical contexts $\hat{\Theta}$ (whose propositions may have existential variables). The definition is straightforward: homomorphically apply the context to the object O , and further define $[\Delta]O$ by induction on Δ .

$$\begin{array}{l}
 [\cdot]O = O \\
 [\Delta, \hat{a} : \tau]O = [\Delta]O \\
 [\Delta, \hat{a} : \tau = t]O = [\Delta]([t/\hat{a}]O)
 \end{array}$$

The order of substitutions in the definition of context application above does not matter because solutions are ground (we may view $[\Delta]O$ as simultaneous substitution). If O only has existential variables from $\text{dom}(\Omega)$, then $[\Omega]O$ is ground.

6.2 Context Extension

The algorithmic context extension judgment $\Theta \vdash \Delta \longrightarrow \Delta'$ says that $\text{dom}(\Delta) = \text{dom}(\Delta')$ and Δ' has the same solutions as Δ , but possibly solves more (that are unsolved in Δ). All typing and subtyping

judgments (under Θ) that have input and output algorithmic contexts Δ and Δ' (respectively) enjoy the property that they increase index information, that is, $\Theta \vdash \Delta \longrightarrow \Delta'$. If $\Theta \vdash \Delta \longrightarrow \Omega$, then Ω *completes* Δ : it has Δ 's solutions, but also solutions to all of Δ 's unsolved variables.

$$\frac{}{\Theta \vdash \cdot \longrightarrow \cdot} \quad \frac{\Theta \vdash \Delta \longrightarrow \Delta'}{\Theta \vdash \Delta, \hat{a} : \tau \longrightarrow \Delta', \hat{a} : \tau} \quad \frac{\Theta \vdash \Delta \longrightarrow \Delta'}{\Theta \vdash \Delta, \hat{a} : \tau=t \longrightarrow \Delta', \hat{a} : \tau=t}$$

$$\frac{\Theta \vdash \Delta \longrightarrow \Delta'}{\Theta \vdash \Delta, \hat{a} : \tau \longrightarrow \Delta', \hat{a} : \tau=t}$$

6.3 Subtyping

Algorithmic subtyping $\Theta; \Delta \vdash A <:^\pm B / W \dashv \Delta'$ says that, under logical context Θ and algorithmic context Δ , the type A is algorithmically a subtype of B if and only if output constraint W holds algorithmically (under suitable solutions including those of Δ'), outputting index solutions Δ' . In subtyping and type equivalence, the delayed output constraints W must remember their logical context via \supset and \forall . For example, in checking that $\exists a : \mathbb{N}. \text{Nat}(a) \wedge (a < 5)$ is a subtype of $\exists a : \mathbb{N}. \text{Nat}(a) \wedge (a < 10)$, the output constraint W is $\forall a : \mathbb{N}. (a < 5) \supset (a < 10)$.

For space reasons, we don't present all algorithmic subtyping rules here (see appendix Fig. 48), but only enough rules to discuss the key design issues. Further, we don't present algorithmic equivalence here (see appendix Figures 44 and 46), which is similar to and simpler than algorithmic subtyping.

In algorithmic subtyping, we maintain the invariant that positive subtypes and negative super-types are ground. The rules

$$\frac{\Theta; \Delta, \hat{a} : \tau \vdash P <:^+ [\hat{a}/a]Q / W \dashv \Delta', \hat{a} : \tau=t}{\Theta; \Delta \vdash P <:^+ \exists a : \tau. Q / W \dashv \Delta'} \quad \frac{\Theta; \Delta, \hat{a} : \tau \vdash [\hat{a}/a]N <:^- M / W \dashv \Delta', \hat{a} : \tau=t}{\Theta; \Delta \vdash \forall a : \tau. N <:^- M / W \dashv \Delta'}$$

are the only subtyping rules which add existential variables (to the side not necessarily ground) to be solved (whereas the declarative system conjures a solution). We pop off the solution as we have the invariant that output contexts are eagerly applied to output constraints and input types.

The rule

$$\frac{t \text{ ground} \quad \Theta; \Delta \vdash F \equiv G / W \dashv \Delta'_1, \hat{a} : \tau, \Delta'_2 \quad \Delta' = \Delta'_1, \hat{a} : \tau=t, \Delta'_2}{\Theta; \Delta \vdash \{v : \mu F \mid (\text{fold}_F \alpha) v =_\tau t\} <:^+ \{v : \mu G \mid (\text{fold}_G \alpha) v =_\tau \hat{a}\} / W \wedge (t = t) \dashv \Delta'} <:^+ / \text{-}\mu\text{INST}$$

runs the functor equivalence algorithm (which outputs constraint W and solutions $\Delta'_1, \hat{a} : \tau, \Delta'_2$), checks that \hat{a} does not get solved there, and then solves \hat{a} to t (yielding Δ') after checking that the latter (which is a subterm of a positive subtype) is ground, outputting the constraint generated by functor equivalence together with the equation $t = t$ (the declarative system can conjure a different but logically equal term for the right-hand side of this equation), and Δ' . Alternatively, there is a rule for when \hat{a} gets solved by functor equivalence, and a rule where a term that is not an existential variable is in place of \hat{a} .

The rule

$$\frac{\Theta; \Delta \vdash M \rightsquigarrow M' [\hat{\Theta}]}{\Theta; \Delta \vdash \downarrow N <:^+ \downarrow M / (\hat{\Theta} \supset^* \underline{N} <:^- M') \dashv \Delta}$$

extracts M' and $\hat{\Theta}$ from M and delays the resulting negative subtyping constraint $\underline{N} <:^- M'$, to be verified under its logical setting $\hat{\Theta}$ (whose propositions, which were extracted from the side not necessarily ground, may have existential variables only solved in value typechecking).

The metaoperation \triangleright^* traverses $\hat{\Theta}$, creating universal quantifiers from universal variables and implications from propositions:

$$\begin{aligned} \cdot \triangleright^* W &= W \\ (\hat{\Theta}, \phi) \triangleright^* W &= \hat{\Theta} \triangleright^* (\phi \triangleright W) \\ (\hat{\Theta}, a : \tau) \triangleright^* W &= \hat{\Theta} \triangleright^* (\forall a : \tau. W) \end{aligned}$$

The dual shift rule is similar. In the declarative system, $\leq^+ \rightsquigarrow L$ and $\leq^- \rightsquigarrow R$ are invertible, which means that they can be eagerly applied without getting stuck; algorithmically, we apply them immediately at polarity shifts, so the above rule corresponds to an algorithmic combination of the declarative rules $\leq^- \rightsquigarrow R$ and $\leq^+ \downarrow$ (and similarly for its dual rule for \uparrow).

For rules with multiple nontrivial premises, such as product subtyping

$$\frac{\Theta; \Delta \vdash P_1 < :^+ Q_1 / W_1 \vdash \Delta'' \quad \Theta; \Delta'' \vdash P_2 < :^+ [\Delta'']Q_2 / W_2 \vdash \Delta'}{\Theta; \Delta \vdash (P_1 \times P_2) < :^+ (Q_1 \times Q_2) / [\Delta']W_1 \wedge W_2 \vdash \Delta'}$$

we thread solutions through inputs, applying them to the non-ground side ($[\Delta]-$ treats Δ as a substitution of index solutions for existential variables), ultimately outputting both delayed constraints. We maintain the invariant that existential variables in output constraints are eagerly solved, which is why, for example, Δ' is applied to W_1 in the conclusion of the above rule, but not to W_2 (that would be redundant).

6.4 Typing

We now discuss issues specific to algorithmic program typing.

Exploiting polarity, we can restrict the flow of index information to the right- and left-focusing stages: in particular, $\Theta; \Delta; \Gamma \vdash v \leftarrow P / \chi \vdash \Delta'$ and $\Theta; \Delta; \Gamma; [N] \vdash s \gg \uparrow P / \chi \vdash \Delta'$, the algorithmic value and spine typechecking judgments. The input types of these judgments can have existential variables, and these judgments synthesize constraints and index solutions, but the algorithmic versions of the other judgments do not; we judgmentally distinguish the latter by replacing the “ \vdash ” in the declarative judgments with “ \triangleright ” (for example, $\Theta; \Gamma \triangleright g \Rightarrow \uparrow P$). Delayed constraints are verified only and immediately after completing a focusing stage, when all their existential variables are necessarily solved.

Consequently, the algorithmic typing judgments for heads, bound expressions, pattern matching, and expressions are essentially the same as their declarative versions, but with a key difference. Namely, in $\text{Alg} \Rightarrow \text{ValAnnot}$, $\text{Alg} \Rightarrow \text{App}$, and $\text{Alg} \leftarrow \uparrow$ (below, respectively), focusing stages start with an empty algorithmic context, outputting ground constraints (and an empty output context because solutions are eagerly applied), and a premise is added to verify these constraints:

$$\frac{\Theta; \Gamma \vdash v \leftarrow P / \chi \vdash \cdot \quad \Theta; \Gamma \triangleleft \chi}{\Theta; \Gamma \triangleright (v : P) \Rightarrow P} \quad \frac{\Theta; \Gamma \triangleright h \Rightarrow \downarrow N \quad \Theta; \Gamma; [N] \vdash s \gg \uparrow P / \chi \vdash \cdot \quad \Theta; \Gamma \triangleleft \chi}{\Theta; \Gamma \triangleright h(s) \Rightarrow \uparrow P}$$

$$\frac{\Theta; \Gamma \vdash v \leftarrow P / \chi \vdash \cdot \quad \Theta; \Gamma \triangleleft \chi}{\Theta; \Gamma \triangleright \text{return } v \leftarrow \uparrow P}$$

Algorithmic typechecking for recursive expressions uses algorithmic subtyping, which outputs a ground constraint W . Because this W is ground, we can verify it ($\Theta \models W$) immediately:

$$\frac{\Theta \vdash N \rightsquigarrow \quad \Theta; \cdot \vdash \forall a : \mathbb{N}. M < :^- N / W \vdash \cdot \quad \Theta \models W}{\Theta; \Gamma \triangleright \text{rec } x : (\forall a : \mathbb{N}. M). e \leftarrow N}$$

For the full definition of algorithmic typing, see appendix Figures 51, 52, and 53.

$$\boxed{\Theta; \Delta; \Gamma \vdash v \Leftarrow P / \chi \vdash \Delta'} \quad \text{Under inputs contexts } \Theta, \Delta, \text{ and } \Gamma, \text{ the value } v \text{ checks against type } P, \\ \text{with output computation constraints } \chi \text{ and output context } \Delta'$$

$$\frac{P \neq \exists, \wedge \quad (x : Q) \in \Gamma \quad \Theta; \Delta \vdash Q < :^+ P / W \vdash \Delta'}{\Theta; \Delta; \Gamma \vdash x \Leftarrow P / W \vdash \Delta'} \text{Alg}\Leftarrow\text{Var}$$

$$\frac{}{\Theta; \Delta; \Gamma \vdash \langle \rangle \Leftarrow 1 / \cdot \vdash \Delta} \text{Alg}\Leftarrow 1$$

$$\frac{\Theta; \Delta; \Gamma \vdash v_1 \Leftarrow P_1 / \chi_1 \vdash \Delta'' \quad \Theta; \Delta''; \Gamma \vdash v_2 \Leftarrow [\Delta'']P_2 / \chi_2 \vdash \Delta'}{\Theta; \Delta; \Gamma \vdash \langle v_1, v_2 \rangle \Leftarrow (P_1 \times P_2) / [\Delta']\chi_1, \chi_2 \vdash \Delta'} \text{Alg}\Leftarrow \times$$

$$\frac{\Theta; \Delta; \Gamma \vdash v_k \Leftarrow P_k / \chi \vdash \Delta'}{\Theta; \Delta; \Gamma \vdash \text{inj}_k v_k \Leftarrow (P_1 + P_2) / \chi \vdash \Delta'} \text{Alg}\Leftarrow +_k$$

$$\frac{\Theta; \Delta, \hat{a} : \tau; \Gamma \vdash v \Leftarrow [\hat{a}/a]P / \chi \vdash \Delta', \hat{a} : \tau = t}{\Theta; \Delta; \Gamma \vdash v \Leftarrow (\exists a : \tau. P) / \chi \vdash \Delta'} \text{Alg}\Leftarrow \exists$$

$$\frac{\Theta; \Delta; \Gamma \vdash v \Leftarrow P / \chi \vdash \Delta'' \quad \Theta; \Delta'' \vdash [\Delta'']\phi \text{ inst} \vdash \Delta'}{\Theta; \Delta; \Gamma \vdash v \Leftarrow (P \wedge \phi) / ([\Delta']\phi, [\Delta']\chi) \vdash \Delta'} \text{Alg}\Leftarrow \wedge$$

$$\frac{; \Theta; \Delta \vdash \{v : F[\mu F] \mid \alpha (F (\text{fold}_F \alpha) v) =_\tau t\} \doteq P \quad \Theta; \Delta; \Gamma \vdash v \Leftarrow P / \chi \vdash \Delta'}{\Theta; \Delta; \Gamma \vdash \text{into}(v) \Leftarrow \{v : \mu F \mid (\text{fold}_F \alpha) v =_\tau t\} / \chi \vdash \Delta'} \text{Alg}\Leftarrow \mu$$

$$\frac{}{\Theta; \Delta; \Gamma \vdash \{e\} \Leftarrow \downarrow N / (e \Leftarrow N) \vdash \Delta} \text{Alg}\Leftarrow \downarrow$$

$$\boxed{\Theta; \Delta; \Gamma; [N] \vdash s \gg \uparrow P / \chi \vdash \Delta'} \quad \text{Under } \Theta, \Delta, \text{ and } \Gamma, \text{ passing spine } s \text{ to a head of type } \downarrow N \\ \text{synthesizes } \uparrow P, \text{ with output constraints } \chi \text{ and context } \Delta'$$

$$\frac{\Theta; \Delta, \hat{a} : \tau; \Gamma; [[\hat{a}/a]N] \vdash s \gg \uparrow P / \chi \vdash \Delta', \hat{a} : \tau = t}{\Theta; \Delta; \Gamma; [\forall a : \tau. N] \vdash s \gg \uparrow P / \chi \vdash \Delta'} \text{AlgSpine}\forall$$

$$\frac{\Theta; \Delta; \Gamma; [N] \vdash s \gg \uparrow P / \chi \vdash \Delta'}{\Theta; \Delta; \Gamma; [\phi \supset N] \vdash s \gg \uparrow P / [\Delta']\phi, \chi \vdash \Delta'} \text{AlgSpine}\supset$$

$$\frac{\Theta; \Delta; \Gamma \vdash v \Leftarrow Q / \chi \vdash \Delta'' \quad \Theta; \Delta''; \Gamma; [[\Delta'']N] \vdash s \gg \uparrow P / \chi' \vdash \Delta'}{\Theta; \Delta; \Gamma; [Q \rightarrow N] \vdash v, s \gg \uparrow P / [\Delta']\chi, \chi' \vdash \Delta'} \text{AlgSpine}\text{App}$$

$$\frac{}{\Theta; \Delta; \Gamma; [\uparrow P] \vdash \cdot \gg \uparrow P / \text{tt} \vdash \Delta} \text{AlgSpine}\text{Nil}$$

Fig. 15. Algorithmic value and spine typing

Besides the instantiation rules (such as $< :^+ / -\mu\text{INST}$) for inductive types in algorithmic subtyping and type equivalence, there are exactly two judgments ($\Theta; \Delta \vdash \phi \text{ inst} \vdash \Delta'$ and $\Theta; \Delta \vdash \phi \equiv \psi \text{ inst} \vdash \Delta'$) responsible for inferring index solutions, both dealing with the output of algorithmic unrolling. Algorithmic unrolling can output indexes of the form $\hat{a} = t$ with t ground, and these equations are solved in either value typechecking, subtyping, or type equivalence. In the former two cases, we can solve \hat{a} as the algebra body t which is necessarily ground (as discussed in Sec. 4.2). The

judgment $\Theta; \Delta \vdash \phi \text{ inst} \dashv \Delta'$ (appendix Fig. 45), used in $\text{Alg} \Leftarrow \wedge$, checks whether ϕ has form $\hat{a} = t$ where t is ground. If so, then it solves \hat{a} to t in Δ ; otherwise, it does not touch Δ .

$$\frac{\Theta \vdash t : \tau}{\Theta; \Delta_1, \hat{a} : \tau, \Delta_2 \vdash \hat{a} = t \text{ inst} \dashv \Delta_1, \hat{a} : \tau=t, \Delta_2} \text{Inst} \quad \frac{\phi \text{ not of form } \hat{a} = t \text{ where } \Theta \vdash t : \tau}{\Theta; \Delta \vdash \phi \text{ inst} \dashv \Delta} \text{NoInst}$$

For example, suppose a head synthesizes $\forall a : \mathbb{N}. \text{Nat}(a) \rightarrow \uparrow \text{Nat}(a)$ and we wish to apply this head to the spine (containing exactly one argument value) $\text{into}(\text{inj}_2 \langle \text{into}(\text{inj}_1 \langle \rangle, \langle \rangle) \rangle)$. We generate a fresh existential variable \hat{a} for a (rule $\text{AlgSpine}\forall$) and then check the value against $\text{Nat}(\hat{a})$ (rule AlgSpineApp). (Checking the same value against type $\exists a : \mathbb{N}. \text{Nat}(a)$ yields the same problem, by dual rule $\text{Alg} \Leftarrow \exists$, and the following solution also works in this case.) The type $\text{Nat}(\hat{a})$ has value-determined index \hat{a} (its Ξ is $\hat{a} : \mathbb{N}$), so it is solvable. We unroll (rule $\text{Alg} \Leftarrow \mu$) $\text{Nat}(\hat{a})$ to $(1 \wedge (\hat{a} = 0)) + (\exists a' : \mathbb{N}. \text{Nat}(a') \times (1 \wedge (\hat{a} = 1 + a')))$ and check $\text{inj}_2 \langle \text{into}(\text{inj}_1 \langle \rangle, \langle \rangle) \rangle$ against that (0 and $1 + a'$ are the bodies of the two branches of Nat 's algebra). In this unrolled type, \hat{a} is no longer tracked by its Ξ , but we can still solve it.

The value now in question is a right injection, so we must check $\langle \text{into}(\text{inj}_1 \langle \rangle, \langle \rangle) \rangle$ against $\exists a' : \mathbb{N}. \text{Nat}(a') \times (1 \wedge (\hat{a} = 1 + a'))$ (rule $\text{Alg} \Leftarrow +_2$). We generate another fresh existential variable \hat{a}' in place of a' . We now check the pair using rule $\text{Alg} \Leftarrow \times$. For the first component, we check $\text{inj}_1 \langle \rangle$ against the unrolled $\text{Nat}(\hat{a}')$, which is $1 \wedge (\hat{a}' = 0) + \dots$. Now we solve $\hat{a}' = 0$ (rules $\text{Alg} \Leftarrow +_1$, $\text{Alg} \Leftarrow \wedge$, Inst , and $\text{Alg} \Leftarrow 1$). This information flows to the type $1 \wedge (\hat{a} = 1 + \hat{a}')$ against which we need to check the second value component $\langle \rangle$. By “this information flows,” we mean that we apply the context output by type checking the first component, namely $\hat{a} : \mathbb{N}, \hat{a}' : \mathbb{N}=0$ (notice \hat{a} is not yet solved), as a substitution to obtain $1 \wedge (\hat{a} = 1 + 0)$ for the second premise of $\text{Alg} \Leftarrow \times$. The right-hand side of the equation now has no existential variables, and we solve $\hat{a} = 1 + 0 = 1$ (again using $\text{Alg} \Leftarrow \wedge$), as expected. It is worth noting that this solving happens entirely within focusing stages.

Values of inductive type may involve program variables, so existential variables may not be solved by $\text{Alg} \Leftarrow \wedge$ (and Inst), but in algorithmic subtyping, using the same instantiation judgment:

$$\frac{\Theta; \Delta \vdash P < :^+ Q / W \dashv \Delta'' \quad \Theta; \Delta'' \vdash [\Delta'']\phi \text{ inst} \dashv \Delta'}{\Theta; \Delta \vdash P < :^+ Q \wedge \phi / [\Delta']W \wedge [\Delta']\phi \dashv \Delta'} < :^+ / \dashv \wedge R$$

Finally, if an equation of the form $\hat{a} = t$ makes its way into type equivalence (by checking a variable value against a sum type), then \hat{a} gets solved, not as t , but rather as the index in the same structural position (including logical structure) of the necessarily ground type on the left of the equivalence (see judgment $\Theta; \Delta \vdash \phi \equiv \psi \text{ inst} \dashv \Delta'$ in appendix Fig. 45, used in appendix Fig. 46).

For example, $b : \mathbb{N}; \hat{a} : \mathbb{N}; x : (1 \wedge (b = 0 + 0)) + \exists b' : \mathbb{N}. \text{Nat}(b') \wedge (b = b' + 0 + 1) \vdash x \Leftarrow P / _ + \hat{a} : \mathbb{N}=b$ where $P = (1 \wedge (\hat{a} = 0)) + \exists a' : \mathbb{N}. \text{Nat}(a') \wedge (\hat{a} = a' + 1)$.

Next, we cover the algorithmic value and spine typing rules (Fig. 15) in detail.

Typechecking values. Because there is no stand-alone algorithmic version of $< :^+ \rightsquigarrow L$ (recall that, in algorithmic subtyping, we eagerly extract immediately under polarity shifts), the rule $\text{Alg} \Leftarrow \text{Var}$ clarifies why we require types in contexts to have already been subjected to extraction. With eager extraction in subtyping under polarity shifts, but without eager type extraction for program variables, we would not be able to extract any assumptions from Q in the (algorithmic) subtyping premise.

Rule $\text{Alg} \Leftarrow \exists$ generates a fresh existential variable which ultimately gets solved within the same stage. Its solution is eagerly applied to the input type and output constraints, so we pop it off of the output context (as it is no longer needed).

Rule $\text{Alg} \Leftarrow \mu$ unrolls the inductive type, checks the inductive type's injected value against the unrolled type, and passes along the constraints and solutions.

Rule $\text{Alg}\Leftarrow\wedge$ delays verifying the validity of the conjoined proposition ϕ until it is grounded. As explained in the example above, existential variables can be solved via propositions generated by algorithmic type unrolling. This is the role of the propositional instantiation judgment used in the second premise: it simply checks whether the proposition is of the form $\hat{a} = t$ where t is ground, in which case it solves \hat{a} as t (rule Inst), and otherwise it does nothing (rule NoInst). If the proposition does solve an existential variable, then the $[\Delta']\phi$ part of the constraint is a trivial equation, but ϕ could be a non-ground proposition unrelated to unrolling, in which case $\Delta' = \Delta''$, whose solutions have not yet been applied to the input ϕ .

Rule $\text{Alg}\Leftarrow\downarrow$ does not have a premise for typechecking the thunked expression (unlike $\text{Decl}\Leftarrow\downarrow$). Instead, the rule delays this typechecking constraint until its existential variables are solved. For example, in

$$;\cdot;\cdot \vdash \langle \{\text{return } \langle \rangle\}, \text{into}(\text{inj}_1 \langle \rangle) \rangle \Leftarrow \exists a : \mathbb{N}. (\downarrow \uparrow (1 \wedge (a = 0))) \times \text{Nat}(a) / \chi \vdash \cdot$$

the output constraint χ has $[0/\hat{a}](\text{return } \langle \rangle \Leftarrow \uparrow (1 \wedge (\hat{a} = 0)))$, where the index solution 0 to the \hat{a} introduced by $\text{Alg}\Leftarrow\exists$ is found only in typechecking the second component of the pair.

Rule $\text{Alg}\Leftarrow 1$ says that $\langle \rangle$ checks against 1, which solves nothing, and there are no further constraints to check.

In rule $\text{Alg}\Leftarrow\times$, we check the first component v_1 , threading through solutions found there in checking the second component v_2 . Checking the second component can solve further existential variables in the first component's constraints χ_1 , but solutions are eagerly applied, so in the conclusion we apply all the solutions only to χ_1 .

Rule $\text{Alg}\Leftarrow+_k$ checks the k -injected value v_k against the sum's k th component type, and passes along the constraints and solutions.

Typechecking spines. Rule $\text{AlgSpine}\forall$, similarly to $\text{Alg}\Leftarrow\exists$, generates a fresh existential variable that ultimately gets solved in typechecking a value (in this case the spine's corresponding value). As usual, we pop off the solution because solutions are eagerly applied.

In rule AlgSpineApp , we typecheck the value v , outputting constraints χ and solutions Δ'' . We thread these solutions through when checking s , the rest of the spine, ultimately outputting constraints χ' and solutions Δ' . The context Δ' may have more solutions than Δ'' , and we eagerly apply solutions, so we need only apply Δ' to the first value's constraints χ .

In $\text{AlgSpine}\supset$, we check the spine s but add the guarding proposition ϕ to the list of constraints to verify later (applying the solutions Δ' found when checking the spine).

In AlgSpineNil , nothing gets solved, so we output the input algorithmic context Δ . Nothing needs to be verified, so we output the trivial constraint tt .

7 ALGORITHMIC METATHEORY

We prove that the algorithmic system (Sec. 6) is decidable, as well as sound and complete with respect to the declarative system (Sec. 4).

7.1 Decidability

We have proved that all algorithmic judgments are decidable (appendix Sec. G). Algorithmic constraint verification $\Theta \models W$ and $\Theta; \Gamma \triangleleft \chi$ boils down to verifying propositional validity $\Theta \vdash \phi$ true, which is known to be decidable [Barrett et al. 2009]. Besides that, our decidability proofs rely on fairly simple metrics for the various algorithmic judgments, which involve a simple size function (appendix Figs. 57 and 58) and counting the number of subtyping constraints W in typing constraint lists χ . We show that, for each algorithmic rule, every premise is smaller than the conclusion, according to the metrics. The most interesting lemmas we use state that the constraints output by

algorithmic equivalence, subtyping and program typing decrease in size (appendix Lemmas G.9, G.10, and G.12). For example:

LEMMA 7.1 (PROGRAM TYPING SHRINKS CONSTRAINTS). *(Lemma G.12 in appendix)*

- (1) If $\Theta; \Delta; \Gamma \vdash v \Leftarrow P / \chi \vdash \Delta'$, then $\text{sz}(\chi) \leq \text{sz}(v)$.
- (2) If $\Theta; \Delta; \Gamma; [N] \vdash s \gg \uparrow P / \chi \vdash \Delta'$, then $\text{sz}(\chi) \leq \text{sz}(s)$.

7.2 Algorithmic Soundness

We show that the algorithmic system is sound with respect to the declarative system. Since the algorithmic system is designed to mimic the judgmental structure of the declarative system, soundness (and completeness) of the algorithmic system is relatively straightforward to prove (the real difficulty lies in designing the overall system to make this the case).

Soundness of algorithmic subtyping says that, if the subtyping algorithm solves indexes under which its verification conditions hold, then subtyping holds declaratively under the same solutions:

THEOREM 7.2 (SOUNDNESS OF ALGORITHMIC SUBTYPING). *(Thm. I.4 in appendix)*

- (1) If $\Theta; \Delta \vdash P <^+ Q / W \vdash \Delta'$ and $\Theta; \Delta \vdash Q \text{ type}[\Xi]$ and P ground and $\Theta \vdash \Delta' \longrightarrow \Omega$ and $\Theta \models [\Omega]W$, then $\Theta \vdash P \leq^+ [\Omega]Q$.
- (2) If $\Theta; \Delta \vdash N <^- M / W \vdash \Delta'$ and $\Theta; \Delta \vdash N \text{ type}[\Xi]$ and M ground and $\Theta \vdash \Delta' \longrightarrow \Omega$ and $\Theta \models [\Omega]W$, then $\Theta \vdash [\Omega]N \leq^- M$.

We prove (appendix Thm. I.4) the soundness of algorithmic subtyping by way of two intermediate (sound) subtyping systems: a declarative system that eagerly extracts under shifts, and a semideclarative system that also eagerly extracts under shifts, but outputs constraints W in the same way as algorithmic subtyping, to be checked by the semideclarative judgment $\Theta \triangleright W$ (that we prove sound with respect to the algorithmic $\Theta \models W$). We define a straightforward subtyping constraint equivalence judgment $\Theta \triangleright W \leftrightarrow W'$, that uses the proposition and type equivalence mentioned in Sec. 4.3, to transport semideclarative to algorithmic subtyping constraint verification (and the other way around for algorithmic subtyping completeness): if $\Theta \triangleright W$ and $\Theta \triangleright W \leftrightarrow W'$, then $\Theta \triangleright W'$ (appendix Lemma H.24).

As a consequence of polarization, the soundness of head, bound expression, expression, and match typing can be stated relatively simply. The typing soundness of values and spines says that if Ω completes the algorithm's solutions such that the algorithm's constraints hold, then the typing of the value or spine holds declaratively with Ω 's solutions.

THEOREM 7.3 (ALG. TYPING SOUND). *(Thm. I.5 in appendix)*

- (1) If $\Theta; \Gamma \triangleright h \Rightarrow P$, then $\Theta; \Gamma \vdash h \Rightarrow P$.
- (2) If $\Theta; \Gamma \triangleright g \Rightarrow \uparrow P$, then $\Theta; \Gamma \vdash g \Rightarrow \uparrow P$.
- (3) If $\Theta; \Delta; \Gamma \vdash v \Leftarrow P / \chi \vdash \Delta'$ and $\Theta; \Delta \vdash P \text{ type}[\Xi]$ and $\Theta \vdash \Delta' \longrightarrow \Omega$ and $\Theta; \Gamma \triangleleft [\Omega]\chi$, then $\Theta; \Gamma \vdash [\Omega]v \Leftarrow [\Omega]P$.
- (4) If $\Theta; \Gamma \triangleright e \Leftarrow N$, then $\Theta; \Gamma \vdash e \Leftarrow N$.
- (5) If $\Theta; \Gamma; [P] \triangleright \{r_i \Rightarrow e_i\}_{i \in I} \Leftarrow N$, then $\Theta; \Gamma; [P] \vdash \{r_i \Rightarrow e_i\}_{i \in I} \Leftarrow N$.
- (6) If $\Theta; \Delta; \Gamma; [N] \vdash s \gg \uparrow P / \chi \vdash \Delta'$ and $\Theta; \Delta \vdash N \text{ type}[\Xi]$ and $\Theta \vdash \Delta' \longrightarrow \Omega$ and $\Theta; \Gamma \triangleleft [\Omega]\chi$, then $\Theta; \Gamma; [[\Omega]N] \vdash [\Omega]s \gg [\Omega]\uparrow P$.

We prove (appendix Thm. I.5) algorithmic typing is sound by a straightforward lexicographic induction on, first, program term structure, and, second, input type size. We happen not to use an intermediate system for this proof, but an intermediate system is very helpful if not indispensable for proving algorithmic typing completeness, discussed next.

7.3 Algorithmic Completeness

We show that the algorithmic system is complete with respect to the declarative system. The declarative system can conjure index solutions that are different from the algorithm's solutions, but they must be equal according to the logical context. We capture this with *relaxed* context extension $\Theta \vdash \Delta \xrightarrow{\sim} \Delta'$, similar to (non-relaxed) context extension ($\Theta \vdash \Delta \longrightarrow \Delta'$) but solutions in Δ may change to equal (under Θ) solutions in Δ' :

$$\frac{\Theta \vdash \Delta \xrightarrow{\sim} \Delta' \quad \Theta \vdash t = t' \text{ true}}{\Theta \vdash \Delta, \hat{a} : \tau=t \xrightarrow{\sim} \Delta', \hat{a} : \tau=t'}$$

Algorithmic completeness says our subtyping algorithm verifies any declarative subtyping. Since declarative subtyping does not eagerly extract from types inside shifts in assumptive positions, but algorithmic subtyping does, the conclusion involves extraction from the given ground type. For example, the equivalence of the algorithmic solutions Δ' to the indexes in Ω for which subtyping declaratively holds may depend on extracted assumptions like Θ_P in part (1) just below.

THEOREM 7.4 (COMPLETENESS OF ALGORITHMIC SUBTYPING). *(Thm. J.14 in appendix)*

- (1) If $\Theta \vdash P \leq^+ [\Omega]Q$ and $\Theta; \Delta \vdash Q \text{ type}[\Xi]$ and P ground and $[\Delta]Q = Q$ and $\Theta \vdash \Delta \xrightarrow{\sim} \Omega$, then there exist P', Θ_P, W , and Δ' such that $\Theta, \Theta_P; \Delta \vdash P' <^+ Q / W \vdash \Delta'$ and $\Theta, \Theta_P \vdash \Delta' \xrightarrow{\sim} \Omega$ and $\Theta, \Theta_P \models [\Omega]W$ and $\Theta \vdash P \rightsquigarrow^+ P' [\Theta_P]$.
- (2) If $\Theta \vdash [\Omega]N \leq^- M$ and $\Theta; \Delta \vdash N \text{ type}[\Xi]$ and M ground and $[\Delta]N = N$ and $\Theta \vdash \Delta \xrightarrow{\sim} \Omega$, then there exist M', Θ_M, W , and Δ' such that $\Theta, \Theta_M; \Delta \vdash N <^- M' / W \vdash \Delta'$ and $\Theta, \Theta_M \vdash \Delta' \xrightarrow{\sim} \Omega$ and $\Theta, \Theta_M \models [\Omega]W$ and $\Theta \vdash M \rightsquigarrow^- M' [\Theta_M]$.

We prove (appendix Thm. J.14) completeness of algorithmic subtyping by using, in a similar way, the same intermediate systems used to prove soundness. However, it's more complicated. Indexes in semideclarative and algorithmic constraints may be syntactically different but logically and semantically equal. More crucially, to prove the completeness of algorithmic typing with respect to semideclarative typing, we need to prove that algorithmic subtyping solves all value-determined indexes of input types that are not necessarily ground:

LEMMA 7.5 (SUB. SOLVES VAL-DET.). *(Lemma J.5 in appendix)*

- (1) If $\Theta; \Delta \vdash P <^+ Q / W \vdash \Delta'$ and P ground and $\Theta; \Delta \vdash Q \text{ type}[\Xi]$, then for all $(\hat{a} : \tau) \in \Xi$, there exists t such that $\Theta \vdash t : \tau$ and $(\hat{a} : \tau=t) \in \Delta'$.
- (2) If $\Theta; \Delta \vdash M <^- N / W \vdash \Delta'$ and N ground and $\Theta; \Delta \vdash M \text{ type}[\Xi]$, then for all $(\hat{a} : \tau) \in \Xi$, there exists t such that $\Theta \vdash t : \tau$ and $(\hat{a} : \tau=t) \in \Delta'$.

We prove (appendix Lemma J.5) this by straightforward induction on the given subtyping derivation, using a similar lemma for type equivalence (appendix Lemma J.4).

We use extraction to achieve a complete subtyping algorithm. For example, the following holds declaratively without extraction but instead using $\leq^+ \wedge L$ (this rule is not in our system; see Sec. 4.5):

$$a : \mathbb{N}, b : \mathbb{N} \vdash \underbrace{[\hat{c} : \mathbb{N}=b](\text{Nat}(\hat{c}) \rightarrow (\hat{c} = b) \supset \uparrow 1)}_{\text{Nat}(b) \rightarrow (b = b) \supset \uparrow 1} \leq^- (\text{Nat}(a) \wedge (a = b)) \rightarrow \uparrow 1$$

However, checking function argument subtyping first, the non-extractive algorithm solves \hat{c} to a (not b) and outputs a verification condition needing $a = b$ to hold under no logical assumptions, which is invalid. Our system instead extracts $a = b$ from the supertype; the algorithmic solution a for \hat{c} and the declarative choice b for \hat{c} are equal under this assumption ($a = b$).

Finally, we prove the completeness of algorithmic typing. Like algorithmic typing soundness, again due to focusing, the head, bound expression, expression, and pattern matching parts are

straightforward to state. But, because algorithmic function application may instantiate indexes different but logically equal to those conjured (semi)declaratively, bound expressions may algorithmically synthesize a type (judgmentally) equivalent to the type it synthesizes declaratively.

We introduced logical context equivalence in Sec. 4.3. Other than in proving that type equivalence implies subtyping, logical context equivalence is used in proving the completeness of algorithmic typing (in particular, we effectively use appendix Lemma B.95 to swap logically equivalent logical contexts in (semi)declarative typing derivations). The type P' in the output of the algorithm in part (6) below can have different index solutions (output Δ') that are logically equal (under Θ) to the solutions in Ω which appear in the declaratively synthesized P . However, P and P' necessarily have the same structure, so $\Theta \vdash P \equiv^+ [\Omega]P'$. Therefore, a bound expression may (semi)declaratively synthesize a type that is judgmentally equivalent to the type synthesized algorithmically. We then extract different but logically equivalent logical contexts from the (equivalent) types synthesized by a bound expression.

As such, algorithmic typing completeness is stated as follows:

THEOREM 7.6 (ALG. TYPING COMPLETE).

(Thm. J.21 in appendix)

- (1) If $\Theta; \Gamma \vdash h \Rightarrow P$, then $\Theta; \Gamma \triangleright h \Rightarrow P$.
- (2) If $\Theta; \Gamma \vdash g \Rightarrow \uparrow P$, then there exists P' such that $\Theta; \Gamma \triangleright g \Rightarrow \uparrow P'$ and $\Theta \vdash P \equiv^+ P'$.
- (3) If $\Theta; \Gamma \vdash v \Leftarrow [\Omega]P$ and $\Theta; \Delta \vdash P$ type $[\Xi]$ and $[\Delta]P = P$ and $\Theta \vdash \Delta \xrightarrow{\sim} \Omega$, then there exist χ and Δ' such that $\Theta; \Delta; \Gamma \vdash v \Leftarrow P / \chi \vdash \Delta'$ and $\Theta \vdash \Delta' \xrightarrow{\sim} \Omega$ and $\Theta; \Gamma \triangleleft [\Omega]\chi$.
- (4) If $\Theta; \Gamma \vdash e \Leftarrow N$, then $\Theta; \Gamma \triangleright e \Leftarrow N$.
- (5) If $\Theta; \Gamma; [P] \vdash \{r_i \Rightarrow e_i\}_{i \in I} \Leftarrow N$, then $\Theta; \Gamma; [P] \triangleright \{r_i \Rightarrow e_i\}_{i \in I} \Leftarrow N$.
- (6) If $\Theta; \Gamma; [[\Omega]N] \vdash s \gg \uparrow P$ and $\Theta; \Delta \vdash N$ type $[\Xi]$ and $[\Delta]N = N$ and $\Theta \vdash \Delta \xrightarrow{\sim} \Omega$, then there exist P' , χ , and Δ' such that $\Theta; \Delta; \Gamma; [N] \vdash s \gg \uparrow P' / \chi \vdash \Delta'$ and $\Theta \vdash \Delta' \xrightarrow{\sim} \Omega$ and $\Theta; \Gamma \triangleleft [\Omega]\chi$ and $\Theta \vdash P \equiv^+ [\Omega]P'$.

We prove (appendix Thm. J.21) algorithmic typing completeness by way of an intermediate, semideclarative typing system, that is essentially the same as declarative typing in that it conjures indexes, but differs in a way similar to algorithmic typing: it outputs constraints χ and only verifies them (via semideclarative $\Theta; \Gamma \widetilde{\triangleleft} \chi$ as opposed to algorithmic $\Theta; \Gamma \triangleleft \chi$) immediately upon completion of focusing stages. Similarly to the proof of algorithmic subtyping completeness, we transport (appendix Lemma J.18) the semideclarative verification of typing constraints over a straightforward typing constraint equivalence judgment $\Theta; \Gamma \triangleleft \chi \leftrightarrow \chi'$ that uses the subtyping constraint equivalence ($\Theta \triangleright W \leftrightarrow W'$) and type equivalence judgments.

To prove that algorithmic typing is complete with respect to semideclarative typing, we use the fact that typing solves all value-determined indexes in input types of focusing stages. This fact is similar to the fact that subtyping solves the value-determined indexes of non-ground types (used in the algorithmic subtyping completeness proof), but the interaction between value-determined indexes and unrolling introduces some complexity: unrolling a refined inductive type does not preserve the type's Ξ . Therefore, we had to split the value typechecking part into the mutually recursive parts (1) and (2); part (3) depends on parts (1) and (2) but not vice versa.

LEMMA 7.7 (TYPING SOLVES VAL-DET.).

(Lemma J.19 in appendix)

- (1) Suppose $\Delta = \Delta_1, \hat{a} : \tau, \Delta_2$. If $\Xi; \Theta; \Delta \vdash \{v : G[\mu F] \mid \beta(G(\text{fold}_F \alpha) v) =_\tau \hat{a}\} \doteq Q$ and $\Theta; \Delta \vdash G$ functor $[\Xi_G]$ and $(\hat{a} : \tau) \notin \Xi_G$ and $\Theta; \Delta; \Gamma \vdash v \Leftarrow Q / \chi \vdash \Delta'$ and $[\Delta]Q = Q$ and $\Theta \vdash \Delta' \xrightarrow{\sim} \Omega$ and $\Theta; \Gamma \widetilde{\triangleleft} [\Omega]\chi$, then there exists t such that $\Theta \vdash t : \tau$ and $(\hat{a} : \tau=t) \in \Delta'$.

- (2) If $\Theta; \Delta; \Gamma \vdash v \Leftarrow P / \chi \dashv \Delta'$
 and $\Theta; \Delta \vdash P \text{ type}[\Xi_P]$ and $[\Delta]P = P$ and $\Theta \vdash \Delta' \xrightarrow{\sim} \Omega$ and $\Theta; \Gamma \preceq [\Omega]\chi$,
 then for all $(\hat{a} : \tau) \in \Xi_P$, there exists t such that $\Theta \vdash t : \tau$ and $(\hat{a} : \tau=t) \in \Delta'$.
- (3) If $\Theta; \Delta; \Gamma; [N] \vdash s \gg \uparrow P / \chi \dashv \Delta'$
 and $\Theta; \Delta \vdash N \text{ type}[\Xi_N]$ and $[\Delta]N = N$ and $\Theta \vdash \Delta' \xrightarrow{\sim} \Omega$ and $\Theta; \Gamma \preceq [\Omega]\chi$,
 then for all $(\hat{a} : \tau) \in \Xi_N$, there exists t such that $\Theta \vdash t : \tau$ and $(\hat{a} : \tau=t) \in \Delta'$.

The proof (appendix Lemma J.19) of part (1) boils down to inversion on the propositional instantiation judgment $\Theta; \Delta \vdash \phi \text{ inst } \dashv \Delta'$ in the unit case of unrolling where ϕ necessarily has the form $\hat{a} = t$ with t ground, due to the invariant that algebras are ground.

8 RELATED WORK

Typing refinement. As far as we know, Constable [1983] was first to introduce the concept of refinement types (though not by that name) as logical subsets of types, writing $\{x : A \mid B\}$ for the *subset type* classifying terms x of type A that satisfy proposition B . Freeman and Pfenning [1991] introduced type refinement to the programming language ML via *datasort* refinements—inclusion hierarchies of ML-style (algebraic, inductive) datatypes—and intersection types for Standard ML: they showed that full type inference is decidable under a refinement restriction, and provided an algorithm based on abstract interpretation. The dangerous interaction of datasort refinements, intersection types, side effects, and call-by-value evaluation was first dealt with by Davies and Pfenning [2000] by way of a value restriction on intersection introduction; they also presented a bidirectional typing algorithm.

Dependent ML (DML) [Xi 1998; Xi and Pfenning 1999] extended the ML type discipline parametrically by *index* domains. DML is only decidable modulo decidability of index constraint satisfiability. DML uses a bidirectional type system with index refinements for a variant of ML, capable of checking properties ranging from in-bound array access [Xi and Pfenning 1998] to program termination [Xi 2002]. DML, similarly to our system, collects constraints from the program and passes them to a constraint solver, but does not guarantee that they are SMT solvable (unlike our system). This is also the approach of systems like Stardust [Dunfield 2007a] (which combines datasort and index refinement, and supports index refinements that are not value-determined, that is, *invaluable refinements*, which we do not consider) and those with liquid types [Rondon et al. 2008]. The latter are based on a Hindley–Milner approach; typically, Hindley–Damas–Milner inference algorithms [Hindley 1969; Milner 1978; Damas and Milner 1982] generate typing constraints to be verified [Heeren et al. 2002].

Due to issues with existential index instantiation, the approach of Xi [1998] (incompletely) translated programs into a let-normal form [Sabry and Felleisen 1993] before typing them, and Dunfield [2007b] provided a complete let-normal translation for similar issues. The system in this paper is already let-normal.

Liquid types. Rondon et al. [2008] introduced *logically qualified data types*, that is, *liquid types*, in a system that extends Hindley–Milner to infer (by abstract interpretation) refinements based on built-in or programmer-provided refinement templates or qualifiers. Kawaguchi et al. [2009] introduced *recursive refinement* via sound and terminating *measures* on algebraic data types; they also introduced *polymorphic refinement*. Vazou et al. [2013] generalize recursive and polymorphic refinement into a single, more expressive mechanism: *abstract refinement*. Our system lacks polymorphism, which we plan to study in future work. Abstract refinements go together with multi-argument measures because abstract refinements may be thought of as predicates of higher-order sort and we can encode multi-argument measures using higher-order sorts. Extending our system with higher-order sorts and abstract refinements is ongoing work. In future work, it would

be interesting to study other features of liquid typing in our setting. Extending our system with liquid inference of refinements, for example, would require adding Hindley–Milner type inference that creates templates, as well as mechanisms to solve these templates, possibly in an initial phase using abstract interpretation.

Unlike DML (and our system), liquid type systems do not distinguish index terms from programs. While this provides simplicity and convenience to the user (from their perspective, there is just one language), it makes it relatively difficult to provide liquid type systems a denotational semantics and to prove soundness results denotationally (rather than operationally), in contrast to our system (we should be able to recover some of this convenience, by requiring users, for example, to make measure annotations like Liquid Haskell). It creates other subtleties such as the fact that annotations for termination metrics in Liquid Haskell must be internally translated to ghost parameters similar to that used in Sec. 3. By contrast, if we extend our system with additional termination metrics, because these metrics are at the index level, we can obviate such ghost parameters. Liquid types’ lack of index distinction also makes it trickier to support computational effects and evaluation orders.

Initial work on liquid types [Rondon et al. 2008; Kawaguchi et al. 2009] used call-by-value languages, but Haskell uses lazy evaluation so Liquid Haskell was discovered to be unsound [Vazou et al. 2014]. Vazou et al. [2014] regained type soundness by imposing operational-semantic restrictions on subtyping and let-binding. In their algorithmic subtyping, there is exactly one rule, \leq -BASE-D, which pertains to refinements of base types (integers, booleans and so on) and inductive data types; however, these types have a *well-formedness restriction*, namely, that the refinement predicates have the type of boolean expressions that reduce to finite values. But this restriction alone does not suffice for soundness under laziness and divergence. As such, their algorithmic typing rule T-CASE-D, which combines let-binding and pattern matching, uses an operational semantics to approximate whether or not the bound expression terminates. If the bound expression might diverge, then so might the entire case expression; otherwise, it checks each branch in a context that assumes the expression reduces to a (potentially infinite Haskell) value.

We also have a type well-formedness restriction, but it is purely syntactic, and only on index quantification, requiring them to be associated with a fold that is necessarily decidable by virtue of a systematic phase distinction between the index level and the program level. Further, via type polarization, our let-binding rule requires the bound expression to return a value, we only allow value types in our program contexts, and we cannot extract index information across polarity shifts (such as in a suspended computation). Therefore, in our system, there is no need to stratify our types according to an approximate criterion; rather, we exploit the systemic distinction between positive (value) types and negative (computation) types, that Levy [2004] designed or discovered to be semantically well-behaved. We suspect that liquid types’ divergence-based stratification is indirectly grappling with logical polarity. Because divergence-based stratification is peculiar to the specific effect of nontermination, it is unclear how their approach may extend to other effects. By way of a standard embedding of CBN or CBV into (our focalized variant of) CBPV we can obtain CBN or CBV subtyping and typing relations automatically respecting any necessary value or covalue restrictions [Zeilberger 2009]. Further, our system is already in a good position to handle the addition of effects other than nontermination.

Contract calculi. Software *contracts* express program properties in the same language as the programs themselves; Findler and Felleisen [2002] introduced contracts for *run-time* verification of higher-order functional programs. These *latent* contracts are not types, but *manifest* contracts are [Greenberg et al. 2010]. Manifest contracts are akin to refinement types. Indeed, Vazou et al. [2013] sketch a proof of type soundness for a liquid type system by translation from liquid types to the

manifest contract calculus F_H of Belo et al. [2011]. However, there is no explicit translation back, from F_H to liquid typing. They mention that the translated terms in F_H do not have *upcasts* because the latter in F_H are logically related to identity functions if they correspond to static subtyping (as they do in the liquid type system): an upcast lemma. Presumably, this facilitates a translation from F_H back to liquid types. However, there are technical problems in F_H that break type soundness and the logical consistency of the F_H contract system; Sekiyama et al. [2017] fix these issues, resulting in the system F_H^σ , but do not consider subtyping and subsumption, and do not prove an upcast lemma.

Bidirectional typing. Bidirectional typing [Pierce and Turner 2000] is a popular way to implement a wide variety of systems, including dependent types [Coquand 1996; Abel et al. 2008], contextual types [Pientka 2008], and object-oriented languages [Odersky et al. 2001]. The bidirectional system of Peyton Jones et al. [2007] supports higher-rank polymorphism. Dunfield and Krishnaswami [2013] also present a bidirectional type system for higher-rank polymorphism, but framed more proof theoretically; Dunfield and Krishnaswami [2019] extend it to a richer language with existentials, indexed types, sums, products, equations over type variables, pattern matching, polarized subtyping, and principality tracking. The bidirectional system of this paper uses logical techniques similar to the systems of Dunfield and Krishnaswami, but it does not consider polymorphism. A survey paper [Dunfield and Krishnaswami 2021] includes some discussion of bidirectional typing’s connections to proof theory. Basically, good bidirectional systems tend to distinguish checking and synthesizing terms or proofs according to their form, such as normal or neutral.

Proof theory, polarization, focusing and analyticity. The concept of polarity most prominent in this paper dates back to Andreoli’s work on focusing for tractable proof search [Andreoli 1992] and Girard’s work on unifying classical, intuitionistic, and linear logic [Girard 1993]. Logical polarity and focusing have been used to explain many common phenomena in programming languages. We mentioned in the overview that Zeilberger [2009] explains the value and evaluation context restrictions in terms of focusing; and Krishnaswami [2009] explains pattern matching as (proof terms of) the left-inversion stage of focused systems (also, that system is bidirectional). More broadly, Downen [2017] discusses many logical dualities common in programming languages.

Brock-Nannestad et al. [2015] study the relation between polarized intuitionistic logic and CBPV. They obtain a bidirectionally typed system of natural deduction related to a variant of the focused sequent calculus *LJF* [Liang and Miller 2009] by η -expansion (for inversion stages). Espírito Santo [2017] does a similar study, but starts with a focused sequent calculus for intuitionistic logic much like the system of Simmons [2014] (but without positive products), proves it equivalent to a natural deduction system (we think the lack of positive products helps establish this equivalence), and defines, also via η -expansion, a variant of CBPV in terms of the natural deduction system. Our system is not in the style of natural deduction, but rather sequent calculus. We think our system relates to CBPV in a similar way—via η -expansion—but we do not prove it in this paper, because we focus on proving type soundness and algorithmic decidability, soundness and completeness.

Barendregt et al. [1983] discovered that a program that typechecks (in a system with intersection types) using subtyping, can also be checked without using subtyping, if the program is sufficiently η -expanded. An analogous phenomenon involving identity coercion was studied by Zeilberger [2009] in a focused setting. Similarly, our ability to place subtyping solely in (value) variable typechecking is achievable due to the focusing (and let-normality) of our system.

Interpreting Kant, Martin-Löf [1994] considers an *analytic* judgment to be one that is derivable using information found only in its inputs (in the sense of the bidirectional modes, input and output). A *synthetic* judgment, in contrast, requires us to look beyond the inputs of the judgment in order to find a derivation. The metatheoretic results for our algorithmic system demonstrate that

our judgments are analytic, *except* the judgment $\Theta \vdash \phi$ true, which is verified by an external SMT solver. As such, our system may be said to be analytic modulo an external SMT solver. Focusing, in proper combination with bidirectional typing (which clarifies where to put type annotations), let-normality and value-determinedness, guarantees that all information needed to generate verification conditions suitable for an SMT solver may be found in the inputs to judgments. In our system, cut formulas can always be inferred from a type annotation or by looking up a variable in the program context, making all our cuts analytic in the sense of Smullyan [1969].

Dependent types. Dependent types, introduced by Martin-Löf [1971, 1975], are a key conceptual and historical precursor to index refinement types. Dependent types may depend on arbitrary program terms, not only terms restricted to indexes. This is highly expressive, but undecidable in languages with divergence. The main difference between refinement and dependent typing is that refinement typing attempts to increase the expressivity of a highly automatic type system, whereas dependent typing attempts to increase the automation of a highly expressive type system. Semantically, refinement type systems differ from dependent type systems in that they refine a pre-existing type system, so that erasure of refinements always preserves typing.

Many dependent type systems impose their own restrictions for the sake of decidability. In Cayenne [Augustsson 1998], typing can only proceed a given number of steps. All well-typed programs in Epigram [McBride and McKinna 2004] are required to terminate so that its type equivalence is decidable. Epigram, and other systems [Chen and Xi 2005; Licata and Harper 2005], allow programmers to write explicit proofs of type equivalence.

Systems like ATS [Xi 2004] and F^{*} [Swamy et al. 2016] can be thought of as combining refinement and dependent types. These systems aim to bring the best of both refinement and dependent types, but ATS is more geared to practical, effectful functional programming (hence refinement types), while F^{*} is more geared to formal verification and dependent types. Unlike our system, they allow the programmer to provide proofs. The overall design of ATS is closer to our system than that of F^{*}, due to its phase distinction between statics and dynamics; but it allows the programmer to write (in the language itself) proofs in order to simplify or eliminate constraints for the (external) constraint solver: Xi calls this *internalized constraint solving*. It should be possible to internalize constraint solving to some extent in our system. Liquid Haskell has a similar mechanism called *refinement reflection* [Vazou et al. 2017] in which programmers can write manual proofs (in Haskell) in cases where automatic Proof by Logical Evaluation and SMT solving fail.

Both ATS and F^{*} have a CBV semantics, which is inherently monadic [Moggi 1989b]. Our system is a variant of CBPV, which subsumes both CBV and CBN. These systems consider effects other than divergence, like exceptions, mutable state and input/output, which we hope to add to our system in future work; this should go relatively smoothly because CBPV is inherently monadic. The system F^{*} allows for termination metrics other than strong induction on naturals, such as lexicographic induction, but we think it would be straightforward to add such metrics to our system, in the way discussed in Sec. 4.7.

Data abstraction and category theory. Categorically, inductive types are initial algebras of endofunctors. We only consider certain *polynomial* endofunctors, which specify tree-shaped or algebraic data structures. Objects (in the sense of object-oriented programming) or coinductive types are dual to inductive types in that, categorically, they are final coalgebras of endofunctors [Cook 2009]. A consideration of categorical duality leads us to a natural (perhaps naïve) question: if we can build a well-behaved system that refines algebraic data types by algebras, could it mean anything to refine objects by coalgebras? We would expect the most direct model of coinductive types would be via negative types, but working out the details is potential future work.

Our rolled refinement types refine *type* constructors μF . Sekiyama et al. [2015], again in work on manifest contracts, compare this to refining (types of) *data* constructors, and provide a translation from type constructor to data constructor refinements. According to Sekiyama et al. [2015] type constructor refinements (such as our $\{v : \mu F \mid (\text{fold}_F \alpha) v =_\tau t\}$) are *easier for the programmer to specify*, but data constructor refinements (such as the output types of our unrolling judgment) are *easier to verify automatically*. Sekiyama et al. [2015] say that their translation from type to data constructor refinements is closely related to the work of Atkey et al. [2012] on refining inductive data in (a fibrational interpretation of) dependently typed languages. Atkey et al. [2012] provide “explicit formulas” computing inductive characterizations of type constructor refinements. These semantic formulas resemble our syntactic unrolling judgment, which may be viewed as a translation from type refinements to data constructor refinements.

Ornaments [McBride 2011] describe how inductive types with different logical or ornamental properties can be systematically related using their algebraic and structural commonalities. Practical work in ornaments seems mostly geared toward code reuse [Dagand and McBride 2012], code refactoring [Williams and Rémy 2017] and such. In contrast, this paper focuses on incorporating similar ideas in a foundational index refinement typing algorithm.

Melliès and Zeilberger [2015] provide a categorical theory of type refinement in general, where functors are considered to be type refinement systems. This framework is based on Reynolds’s distinction between *intrinsic* (or Church) and *extrinsic* (or Curry) views of typing [Reynolds 1998]. We think that our system fits into this framework, but haven’t confirmed it formally. This is most readily seen in the fact that the semantics of our refined system is simply the semantics (*intrinsic* to unrefined typing derivations) of its erasure of indexes, which express *extrinsic* properties of (erased) programs.

9 CONCLUSION AND FUTURE WORK

We have presented a declarative system for index-based recursive refinement typing (with nullary measures) that is logically designed, semantically sound, and theoretically implementable. We have proved that our declarative system is sound with respect to an elementary domain-theoretic denotational semantics, which implies that our system is logically consistent and totally correct. We have also presented an algorithmic system and proved it is decidable, as well as sound and complete with respect to the declarative system. Focusing yields CBPV, which already has a clear denotational semantics, and refining it by an index domain (and by measures in it) facilitates a semantics in line with the perspective of Melliès and Zeilberger [2015]. But focusing (in combination with value-determinedness) also allows for an easy proof of the completeness of a decidable typing algorithm. The relative ease with which we demonstrate (denotational-semantic) soundness and (completeness of) decidable typing for a realistic language follows from a single, proof-theoretic technique: focusing.

Researchers of liquid typing have laid out an impressive and extensive research program providing many useful features which would be very interesting to study in our setting. We plan to add parametric polymorphism in future work, which goes along with refinement abstraction [Vazou et al. 2013]. Refinement abstraction may be thought of as predicates of higher-order sort, which can also accommodate multi-argument measures (such as whether a list of naturals is in increasing order). We are adding multi-argument measures and refinement abstraction in ongoing work. It is also of great interest to study other features of liquid typing, like refinement inference with templates and refinement reflection, though arguably the latter is more closely related to dependent typing.

We also plan to allow the use of multiple measures on inductive types (so we can specify, for example, the type of length- n lists of naturals in increasing order). It would also be interesting to

experiment with our value-determinedness technique. By allowing quantification over indexes in propositions whose only variable dependencies are *value-determined*, we think we can simulate termination metrics by other ones (such as $<$ on *sums* of natural numbers by $<$ on natural numbers as such).

In future work, we hope to apply our type refinement system (or future extensions of it) to various domains, from static time complexity analysis [Wang et al. 2017] to resource analysis [Handley et al. 2019]. Eventually, we hope to be able to express, for example, that a program terminates within a worst-case amount of time. Our system is parametric in the index domain, provided it satisfies some basic properties. Different index domains may be suitable for different applications. We also hope to add more effects, such as input/output and mutable reference cells. CBPV is built for effects, but our refinement layer may result in interesting interactions between effects and indexes.

Our system may at first seem complicated, but its metatheoretic proofs are largely straightforward, if lengthy (at least as presented). A primary source of this complexity is the proliferation of judgments. However, having various judgments helps us organize different forms of knowledge [Martin-Löf 1996] or (from a Curry–Howard perspective) stages or parts of an implementation (such as pattern-matching, processing an argument list, and so on).

Our system focuses on the feature of nullary measures of algebraic data types, and does not include key typing features expected of a realistic programming language (such as additional effects and polymorphism, which we hope to add in future work). Adding such expressive features tends to significantly affect the metatheory and the techniques used to prove it. We hope to reflect on the development of our proofs (including those for systems with polymorphism [Dunfield and Krishnaswami 2013]) in search of abstractions which may help designers of practical, general-purpose functional languages to establish crucial metatheoretic properties.

ACKNOWLEDGMENTS

We thank the anonymous reviewers for their thorough reading and recommendations, which helped to improve our paper. We also thank Ondrej Baranovič [2023] for implementing the system presented in this paper. This work was supported in part by the Natural Sciences and Engineering Research Council of Canada through Discovery Grant RGPIN–2018–04352, and also in part by European Research Council (ERC) Consolidator Grant for the project “TypeFoundry”, funded under the European Union’s Horizon 2020 Framework Programme (grant agreement no. 101002277).

REFERENCES

- Andreas Abel, Thierry Coquand, and Peter Dybjer. 2008. Verifying a Semantic $\beta\eta$ -Conversion Test for Martin-Löf Type Theory. In *Mathematics of Program Construction (MPC’08) (LNCS)*, Vol. 5133. Springer, 29–56.
- Jean-Marc Andreoli. 1992. Logic programming with focusing proofs in linear logic. *J. Logic and Computation* 2, 3 (1992), 297–347.
- Robert Atkey, Patricia Johann, and Neil Ghani. 2012. Refining Inductive Types. *Logical Methods in Computer Science* Volume 8, Issue 2 (June 2012). [https://doi.org/10.2168/LMCS-8\(2:9\)2012](https://doi.org/10.2168/LMCS-8(2:9)2012)
- Lennart Augustsson. 1998. Cayenne—a Language with Dependent Types. In *ICFP*. 239–250.
- Ondrej Baranovič. 2023. LTR. (2023). <https://github.com/nulano/LTR>.
- Henk Barendregt, Mario Coppo, and Mariangiola Dezani-Ciancaglini. 1983. A Filter Lambda Model and the Completeness of Type Assignment. *J. Symbolic Logic* 48, 4 (1983), 931–940.
- Clark Barrett, Roberto Sebastiani, Sanjit A. Seshia, and Cesare Tinelli. 2009. *Satisfiability modulo theories* (1 ed.). Number 1 in Frontiers in Artificial Intelligence and Applications. IOS Press, 825–885. <https://doi.org/10.3233/978-1-58603-929-5-825>
- João Filipe Belo, Michael Greenberg, Atsushi Igarashi, and Benjamin C. Pierce. 2011. Polymorphic Contracts. In *Proceedings of the 20th European Symposium on Programming (Lecture Notes in Computer Science)*, Gilles Barthe (Ed.), Vol. 6602. Springer International Publishing, 18–37. https://doi.org/10.1007/978-3-642-19718-5_2
- Taus Brock-Nannestad, Nicolas Guenot, and Daniel Gustafsson. 2015. Computation in Focused Intuitionistic Logic. In *17th International Symposium on Principles and Practice of Declarative Programming (PPDP ’15)*. ACM Press, New York, NY, USA, 43–54. <https://doi.org/10.1145/2790449.2790528>

- Iliano Cervesato and Frank Pfenning. 2003. A Linear Spine Calculus. *J. Logic and Computation* 13, 5 (2003), 639–688.
- Chiyan Chen and Hongwei Xi. 2005. Combining programming with theorem proving. In *ICFP*. 66–77.
- James Cheney and Ralf Hinze. 2003. *First-class phantom types*. Technical Report CUCIS TR2003-1901. Cornell University. <https://hdl.handle.net/1813/5614>
- Robert L. Constable. 1983. Mathematics as Programming. In *Logics of Programs, Workshop, Carnegie Mellon University, Pittsburgh, PA, USA, June 6-8, 1983, Proceedings (Lecture Notes in Computer Science)*, Edmund M. Clarke and Dexter Kozen (Eds.), Vol. 164. Springer, 116–128. https://doi.org/10.1007/3-540-12896-4_359
- William R. Cook. 2009. On Understanding Data Abstraction, Revisited. In *Proceedings of the 24th ACM SIGPLAN Conference on Object Oriented Programming Systems Languages and Applications (OOPSLA '09)*. ACM Press, New York, NY, USA, 557–572. <https://doi.org/10.1145/1640089.1640133>
- Thierry Coquand. 1996. An Algorithm for Type-Checking Dependent Types. *Science of Computer Programming* 26, 1–3 (1996), 167–177.
- Pierre-Evariste Dagand and Conor McBride. 2012. Transporting Functions across Ornaments. In *Proceedings of the 17th ACM SIGPLAN International Conference on Functional Programming (ICFP '12)*. ACM Press, New York, NY, USA, 103–114. <https://doi.org/10.1145/2364527.2364544>
- Luis Damas and Robin Milner. 1982. Principal type-schemes for functional programs. In *POPL*. ACM Press, 207–212.
- Rowan Davies and Frank Pfenning. 2000. Intersection Types and Computational Effects. In *ICFP*. ACM Press, 198–208.
- Leonardo de Moura and Nikolaj Bjørner. 2008. Z3: An Efficient SMT Solver. In *TACAS (Lecture Notes in Computer Science)*, Vol. 4963. Springer, 337–340.
- Paul Downen. 2017. *Sequent Calculus: A Logic and a Language for Computation and Duality*. Ph.D. Dissertation. University of Oregon. <https://www.cs.uoregon.edu/Reports/PHD-201706-Downen.pdf>.
- Jana Dunfield. 2007a. Refined typechecking with Stardust. In *Programming Languages meets Programming Verification (PLPV '07)*. ACM Press, 21–32.
- Jana Dunfield. 2007b. *A Unified System of Type Refinements*. Ph.D. Dissertation. Carnegie Mellon University. CMU-CS-07-129.
- Jana Dunfield and Neelakantan R. Krishnaswami. 2013. Complete and Easy Bidirectional Typechecking for Higher-Rank Polymorphism. In *ICFP*. ACM Press, 429–442. arXiv:1306.6032 [cs.PL].
- Jana Dunfield and Neelakantan R. Krishnaswami. 2019. Sound and complete bidirectional typechecking for higher-rank polymorphism with existentials and indexed types. *Proc. ACM Program. Lang.* 3, POPL (2019), 9:1–9:28. <https://doi.org/10.1145/3290322>
- Jana Dunfield and Neelakantan R. Krishnaswami. 2021. Bidirectional Typing. *ACM Comput. Surv.* 54, 5 (2021). <https://doi.org/10.1145/3450952>
- Jana Dunfield and Frank Pfenning. 2003. Type Assignment for Intersections and Unions in Call-by-Value Languages. In *FoSSaCS*. Springer, 250–266.
- José Espírito Santo. 2017. The Polarized λ -calculus. *Electronic Notes in Theoretical Computer Science* 332 (2017), 149–168. <https://doi.org/10.1016/j.entcs.2017.04.010> LSF A 2016 - 11th Workshop on Logical and Semantic Frameworks with Applications (LSFA).
- Robert Bruce Findler and Matthias Felleisen. 2002. Contracts for Higher-Order Functions. In *ICFP*. 48–59.
- Cormac Flanagan, Amr Sabry, Bruce F. Duba, and Matthias Felleisen. 1993. The essence of compiling with continuations. In *PLDI*. 237–247.
- Robert W. Floyd. 1967. Assigning Meanings to Programs. *Proceedings of Symposium on Applied Mathematics* 19 (1967), 19–32.
- Tim Freeman and Frank Pfenning. 1991. Refinement Types for ML. In *PLDI*. ACM Press, 268–277.
- Jean-Yves Girard. 1993. On the Unity of Logic. *Annals of Pure and Applied Logic* 59, 3 (1993), 201–217. [https://doi.org/10.1016/0168-0072\(93\)90093-s](https://doi.org/10.1016/0168-0072(93)90093-s)
- Jean-Yves Girard. 1992. A Fixpoint Theorem in Linear Logic. (1992). Post to Linear Logic mailing list, <http://www.seas.upenn.edu/~sweirich/types/archive/1992/msg00030.html>.
- Michael Greenberg, Benjamin C. Pierce, and Stephanie Weirich. 2010. Contracts Made Manifest. In *Proceedings of the 37th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '10)*. ACM Press, New York, NY, USA, 353–364. <https://doi.org/10.1145/1706299.1706341>
- Carl A. Gunter. 1993. *Semantics of programming languages - structures and techniques*. MIT Press.
- Martin A. T. Handley, Niki Vazou, and Graham Hutton. 2019. Liquidate Your Assets: Reasoning about Resource Usage in Liquid Haskell. *Proc. ACM Program. Lang.* 4, POPL, Article 24 (dec 2019), 27 pages. <https://doi.org/10.1145/3371092>
- Bob Harper and Mark Lillibridge. 1991. ML with callcc is unsound. Post to TYPES mailing list, 8 July 1991, archived at <https://www.seas.upenn.edu/~sweirich/types/archive/1991/msg00034.html>. (1991).
- Robert Harper, John C. Mitchell, and Eugenio Moggi. 1990. Higher-Order Modules and the Phase Distinction. In *Proceedings of the 17th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. ACM Press, San Francisco, California, USA, 341–354. <https://doi.org/10.1145/96709.96744>

- Bastiaan Heeren, Jurriaan Hage, and S. Doaitse Swierstra. 2002. *Generalizing Hindley-Milner Type Inference Algorithms*. Technical Report UU-CS-2002-031. Department of Information and Computing Sciences, Utrecht University. <http://www.cs.uu.nl/research/techreps/repo/CS-2002/2002-031.pdf>
- R. Hindley. 1969. The principal type-scheme of an object in combinatory logic. *Trans. Amer. Math. Soc.* 146 (1969), 29–60.
- C. A. R. Hoare. 1969. An Axiomatic Basis for Computer Programming. *Commun. ACM* 12, 10 (Oct. 1969), 576–580. <https://doi.org/10.1145/363235.363259>
- Ming Kawaguchi, Patrick Rondon, and Ranjit Jhala. 2009. Type-Based Data Structure Verification. In *Proceedings of the 30th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI '09)*. ACM Press, New York, NY, USA, 304–315. <https://doi.org/10.1145/1542476.1542510>
- Andrew Kennedy. 1994. Dimension Types. In *European Symposium on Programming (ESOP '94)*, Vol. 788. Springer, 348–362.
- Neelakantan R. Krishnaswami. 2009. Focusing on Pattern Matching. In *POPL*. ACM Press, 366–378.
- Nico Lehmann, Adam T. Geller, Niki Vazou, and Ranjit Jhala. 2023. Flux: Liquid Types for Rust. *Proc. ACM Program. Lang.* 7, PLDI, Article 169 (jun 2023), 25 pages. <https://doi.org/10.1145/3591283>
- Paul Blain Levy. 2004. *Call-By-Push-Value: A Functional/Imperative Synthesis (Semantics Structures in Computation, V. 2)*. Kluwer Academic Publishers, Norwell, MA, US.
- Chuck Liang and Dale Miller. 2009. Focusing and polarization in linear, intuitionistic, and classical logics. *Theoretical Computer Science* 410, 46 (2009), 4747–4768. <https://doi.org/10.1016/j.tcs.2009.07.041> Abstract Interpretation and Logic Programming: In honor of professor Giorgio Levi.
- Daniel R. Licata and Robert Harper. 2005. *A formulation of Dependent ML with explicit equality proofs*. Technical Report CMU-CS-05-178. Carnegie Mellon University. <https://doi.org/10.1184/R1/6587429.v1>
- Per Martin-Löf. 1971. A Theory of Types. (1971). Manuscript, Stockholm University. <https://raw.githubusercontent.com/michaelt/martin-lof/master/pdfs/martin-loef1971%20-%20A%20Theory%20of%20Types.pdf>
- Per Martin-Löf. 1975. An Intuitionistic Theory of Types: Predicative Part. In *Logic Colloquium '73*, H. E. Rose and J. C. Shepherdson (Eds.). Studies in Logic and the Foundations of Mathematics, Vol. 80. Elsevier, 73–118. [https://doi.org/10.1016/S0049-237X\(08\)71945-1](https://doi.org/10.1016/S0049-237X(08)71945-1)
- Per Martin-Löf. 1984. *Intuitionistic type theory*. Studies in Proof Theory, Vol. 1. Bibliopolis. iv+91 pages.
- Per Martin-Löf. 1994. Analytic and Synthetic Judgements in Type Theory. In *Kant and Contemporary Epistemology*, Paolo Parrini (Ed.). Springer Netherlands, 87–99.
- Per Martin-Löf. 1996. On the meanings of the logical constants and the justifications of the logical laws. *Nordic Journal of Philosophical Logic* 1, 1 (1996), 11–60.
- Conor McBride. 2011. Ornamental Algebras, Algebraic Ornaments. <https://personal.cis.strath.ac.uk/conor.mcbride/pub/OAAO/LitOrn.pdf>
- Conor McBride and James McKinna. 2004. The view from the left. *J. Functional Programming* 14, 1 (2004), 69–111.
- Paul-André Mellès and Noam Zeilberger. 2015. Functors Are Type Refinement Systems. In *Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '15)*. ACM Press, New York, NY, USA, 3–16. <https://doi.org/10.1145/2676726.2676970>
- Robin Milner. 1978. A theory of type polymorphism in programming. *J. Computer and System Sciences* 17, 3 (1978), 348–375.
- Eugenio Moggi. 1989a. A Category-Theoretic Account of Program Modules. In *Category Theory and Computer Science*. Springer-Verlag, Berlin, Heidelberg, 101–117.
- Eugenio Moggi. 1989b. Computational Lambda-Calculus and Monads. In *Logic in Computer Science (LICS '89)*. 14–23.
- Martin Odersky, Matthias Zenger, and Christoph Zenger. 2001. Colored Local Type Inference. In *POPL*. ACM Press, 41–53.
- Simon Peyton Jones, Dimitrios Vytiniotis, Stephanie Weirich, and Mark Shields. 2007. Practical type inference for arbitrary-rank types. *J. Functional Programming* 17, 1 (2007), 1–82.
- Frank Pfenning. 2008. Church and Curry: Combining intrinsic and extrinsic typing. In *Reasoning in Simple Type Theory: Festschrift in Honor of Peter B. Andrews on His 70th Birthday*. College Publications. <http://www.cs.cmu.edu/~fp/papers/andrews08.pdf>
- Brigitte Pientka. 2008. A type-theoretic foundation for programming with higher-order abstract syntax and first-class substitutions. In *POPL*. ACM Press, 371–382.
- Benjamin C. Pierce and David N. Turner. 2000. Local Type Inference. *ACM Trans. Prog. Lang. Sys.* 22 (2000), 1–44.
- John C. Reynolds. 1998. *Theories of programming languages*. Cambridge University Press.
- Nick Rioux and Steve Zdancewic. 2020. Computation Focusing. *Proc. ACM Program. Lang.* 4, ICFP, Article 95 (aug 2020), 27 pages. <https://doi.org/10.1145/3408977>
- Patrick Rondon, Ming Kawaguchi, and Ranjit Jhala. 2008. Liquid types. In *PLDI*. ACM Press, 159–169.
- Amr Sabry and Matthias Felleisen. 1993. Reasoning about Programs in Continuation-Passing Style. *Lisp Symb. Comput.* 6, 3–4 (Nov. 1993), 289–360. <https://doi.org/10.1007/BF01019462>
- Peter Schroeder-Heister. 1994. Definitional reflection and the completion. In *Extensions of Logic Programming (LNCS)*. Springer, 333–347.

- Taro Sekiyama, Atsushi Igarashi, and Michael Greenberg. 2017. Polymorphic Manifest Contracts, Revised and Resolved. *ACM Trans. Program. Lang. Syst.* 39, 1, Article 3 (Feb. 2017), 36 pages. <https://doi.org/10.1145/2994594>
- Taro Sekiyama, Yuki Nishida, and Atsushi Igarashi. 2015. Manifest Contracts for Datatypes. In *Proceedings of the 42nd Symposium on Principles of Programming Languages*. ACM, 195–207. <https://doi.org/10.1145/2676726.2676996>
- Robert J. Simmons. 2014. Structural Focalization. *ACM Trans. Comput. Logic* 15, 3, Article 21 (Sept. 2014). <https://doi.org/10.1145/2629678>
- Raymond M. Smullyan. 1969. Analytic cut. *Journal of Symbolic Logic* 33 (1969), 560 – 564.
- Nikhil Swamy, Cătălin Hrițcu, Chantal Keller, Aseem Rastogi, Antoine Delignat-Lavaud, Simon Forest, Karthikeyan Bhargavan, Cédric Fournet, Pierre-Yves Strub, Markulf Kohlweiss, Jean-Karim Zinzindohoue, and Santiago Zanella-Béguelin. 2016. Dependent Types and Multi-Monadic Effects in F*. In *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '16)*. ACM Press, New York, NY, USA, 256–270. <https://doi.org/10.1145/2837614.2837655>
- W. W. Tait. 1967. Intensional Interpretations of Functionals of Finite Type I. *J. Symbolic Logic* 32, 2 (1967), 198–212. <http://www.jstor.org/stable/2271658>
- Alan M. Turing. 1936. On Computable Numbers, with an Application to the Entscheidungsproblem. *Proceedings of the London Mathematical Society* 2, 42 (1936), 230–265.
- Niki Vazou, Patrick Maxim Rondon, and Ranjit Jhala. 2013. Abstract Refinement Types. In *Programming Languages and Systems - 22nd European Symposium on Programming, ESOP 2013, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2013, Rome, Italy, March 16-24, 2013. Proceedings (Lecture Notes in Computer Science)*, Matthias Felleisen and Philippa Gardner (Eds.), Vol. 7792. Springer, 209–228. https://doi.org/10.1007/978-3-642-37036-6_13
- Niki Vazou, Eric L. Seidel, Ranjit Jhala, Dimitrios Vytiniotis, and Simon Peyton-Jones. 2014. Refinement Types for Haskell. In *Proceedings of the 19th ACM SIGPLAN International Conference on Functional Programming (ICFP '14)*. ACM Press, New York, NY, USA, 269–282. <https://doi.org/10.1145/2628136.2628161>
- Niki Vazou, Anish Tondwalkar, Vikraman Choudhury, Ryan G. Scott, Ryan R. Newton, Philip Wadler, and Ranjit Jhala. 2017. Refinement Reflection: Complete Verification with SMT. *Proc. ACM Program. Lang.* 2, POPL, Article 53 (dec 2017), 31 pages. <https://doi.org/10.1145/3158141>
- Peng Wang, Di Wang, and Adam Chlipala. 2017. TiML: A Functional Language for Practical Complexity Analysis with Invariants. *Proc. ACM Program. Lang.* 1, OOPSLA, Article 79 (Oct. 2017), 26 pages. <https://doi.org/10.1145/3133903>
- Kevin Watkins, Iliano Cervesato, Frank Pfenning, and David Walker. 2004. A Concurrent Logical Framework: The Propositional Fragment. In *Types for Proofs and Programs (TYPES 2003)*. Springer LNCS 3085, 355–377.
- Thomas Williams and Didier Rémy. 2017. A Principled Approach to Ornamentation in ML. *Proc. ACM Program. Lang.* 2, POPL, Article 21 (Dec. 2017). <https://doi.org/10.1145/3158109>
- Andrew K. Wright. 1995. Simple Imperative Polymorphism. *Lisp and Symbolic Computation* 8, 4 (1995), 343–355.
- Hongwei Xi. 1998. *Dependent Types in Practical Programming*. Ph.D. Dissertation. Carnegie Mellon University. <https://www.cs.cmu.edu/~rwh/students/xi.pdf>
- Hongwei Xi. 2002. Dependent Types for Program Termination Verification. *Journal of Higher-Order and Symbolic Computation* 15 (Oct. 2002), 91–131.
- Hongwei Xi. 2004. Applied Type System (extended abstract). In *TYPES 2003 (LNCS)*. Springer, 394–408.
- Hongwei Xi, Chiyan Chen, and Gang Chen. 2003. Guarded recursive datatype constructors. In *POPL*. ACM Press, 224–235.
- Hongwei Xi and Frank Pfenning. 1998. Eliminating array bound checking through dependent types. In *PLDI*. 249–257.
- Hongwei Xi and Frank Pfenning. 1999. Dependent Types in Practical Programming. In *POPL*. ACM Press, 214–227.
- Noam Zeilberger. 2009. Refinement Types and Computational Duality. In *Programming Languages meets Programming Verification (PLPV '09)*. ACM Press, 15–26.