



Polymorphic Type Inference for Dynamic Languages

GIUSEPPE CASTAGNA, CNRS, Université Paris Cité, France

MICKAËL LAURENT, Université Paris Cité, France

KIM NGUYỄN, Université Paris-Saclay, France

We present a type system that combines, in a controlled way, first-order polymorphism with intersection types, union types, and subtyping, and prove its safety. We then define a type reconstruction algorithm that is sound and terminating. This yields a system in which unannotated functions are given polymorphic types (thanks to Hindley-Milner) that can express the overloaded behavior of the functions they type (thanks to the intersection introduction rule) and that are deduced by applying advanced techniques of type narrowing (thanks to the union elimination rule). This makes the system a prime candidate to type dynamic languages.

CCS Concepts: • **Theory of computation** → **Type structures; Program analysis**; • **Software and its engineering** → **Functional languages; Polymorphism**.

Additional Key Words and Phrases: polymorphism, union types, intersection types, type reconstruction.

ACM Reference Format:

Giuseppe Castagna, Mickaël Laurent, and Kim Nguyễn. 2024. Polymorphic Type Inference for Dynamic Languages. *Proc. ACM Program. Lang.* 8, POPL, Article 40 (January 2024), 32 pages. <https://doi.org/10.1145/3632882>

1 INTRODUCTION

Typing dynamic languages is a challenging endeavour even for very simple pieces of code. For instance, JavaScript’s logical or operator “`||`” behaves like the following function (also in JavaScript):¹

```
1 function lor (x, y) {  
2     if (x) { return x; } else { return y; }  
3 }
```

A naive type for this function is $(\text{Bool}, \text{Bool}) \rightarrow \text{Bool}$, which states that `lor` is a function that takes two Boolean arguments and returns a Boolean result. This however is an overly restrictive type, that does not account for the fact that in JavaScript logical operators such as `lor` can be applied to any pairs of arguments, not just to Boolean ones. JavaScript distinguishes two kinds of values: eight “falsy” values (i.e., `false`, `""`, `0`, `-0`, `0n`, `undefined`, `null`, and `NaN`) and the “truthy” values (all the others). The expression `if` executes the `else` code if and only if the tested value is falsy. If we want to change the previous type to account for this fact, then we should give `lor` the type $(\text{Any}, \text{Any}) \rightarrow \text{Any}$ (where `Any` is the type of all values), which is a rather useless type since it essentially states that `lor` is a binary function. To give `lor` a more informative type, we need union and intersection types (which are already integrated in typed versions of JavaScript such as

¹This definition does not capture the short-circuit evaluation of “`||`”.

Authors’ addresses: Giuseppe Castagna, Mickaël Laurent, Institut de Recherche en Informatique Fondamentale (IRIF), Université Paris Cité, CNRS, 8 place Aurélie Nemours, 75013 Paris, France; Kim Nguyen, Laboratoire Méthodes Formelles, Université Paris-Saclay, CNRS, ENS Paris-Saclay, 91190, Gif-sur-Yvette, France.



This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2024 Copyright held by the owner/author(s).

ACM 2475-1421/2024/1-ART40

<https://doi.org/10.1145/3632882>

TypeScript [Microsoft] and Flow [Facebook]): we define the type Falsy as the following union type `false` \vee `""` \vee `0` \vee `-0` \vee `0n` \vee `undefined` \vee `null` \vee `NaN`, where each value denotes here the *singleton type* containing that value, and the type Truthy to be its complement, \neg Falsy, that is, the type of all values that are not of type Falsy. Then we can deduce for `10r` the following more precise type

$$((\text{Truthy}, \text{Any}) \rightarrow \text{Truthy}) \wedge ((\text{Falsy}, \text{Truthy}) \rightarrow \text{Truthy}) \wedge ((\text{Falsy}, \text{Falsy}) \rightarrow \text{Falsy}) \quad (1)$$

In this type, \wedge is a type combinator denoting intersection and meaning that the function has all the types given in the intersection: that is, in words, if the first argument of a function of this type is a Truthy, then the function returns a Truthy regardless of the second argument (first arrow type), while if the first argument is a Falsy, then the result is of the same type as the second argument's type (second and third arrow type). Notice how the use of an intersection of arrow types corresponds to the typing of an “overloading” behavior (also known as, *ad hoc* polymorphism [Strachey 1967]), insofar as the result of an application depends on the *type* of the input.

In order to derive such a type, the type system must deduce that whenever the condition tested by the `if` holds, then `x` is of type Truthy and, therefore, (i) that all occurrences of `x` in the “then” branch (here just one) have type Truthy and (ii) that all the occurrences of the same variable `x` in the branch “else” (here none) have thus type Falsy. This kind of deduction is usually referred as *type narrowing* or *occurrence typing* since it requires to “narrow” the type of a variable `x` differently for its different occurrences. A type system such as the one for Typed Racket—defined in [Tobin-Hochstadt and Felleisen 2010] where the term *occurrence typing* was first introduced—is able to *check* that `10r` has the type in (1), meaning that the deduction requires the programmer to explicitly specify the type in the code. The system by Castagna et al. [2022b] makes a step further, since not only it can check that `10r` has the type in (1), but also it can reconstruct for `10r` the intersection type $((\text{Truthy}, \text{Any}) \rightarrow \text{Truthy}) \wedge ((\text{Falsy}, \text{Any}) \rightarrow \text{Any})$ which, although it is less precise a type than (1), it is inferred from the code of `10r` as is, without needing any type annotation. This latter work constitutes the state of the art of this kind of inference, since it is the only system that can reconstruct intersections of arrow types.

In this work we go a step further, and show how to infer (i.e., reconstruct) intersections of *polymorphic* function types. In particular, the system we present here reconstructs for `10r` the following first order polymorphic type (where α and β are type variables):²

$$\forall \alpha, \beta. ((\alpha \wedge \text{Truthy}, \text{Any}) \rightarrow \alpha \wedge \text{Truthy}) \wedge ((\text{Falsy}, \beta) \rightarrow \beta) \quad (2)$$

This type completely specifies the semantics of the function `10r`: it states that if the first argument is a Truthy, then the application of the function returns the first argument,³ otherwise it returns the second argument. This type is more precise than the one in (1), since it allows the system to deduce that, say, if the first argument of `10r` is an object, then the result will be an object of the same type (rather than just a truthy value). Not only does the system we present here infers such a precise type, but this kind of precision is compositional, yielding an accurate type also for the expressions in which the function is used. For instance, if we define the following function:

```

4  function id (x) {
5      return 10r(x, x)
6  }
```

²This type can be considered as an encoding of $\forall (\alpha \leq \text{Truthy}). \forall (\beta). ((\alpha, \text{Any}) \rightarrow \alpha) \wedge ((\text{Falsy}, \beta) \rightarrow \beta)$ a type expressed in so-called bounded polymorphism: see Castagna [2023a, Section 2].

³Strictly speaking, the type states that the function returns a result of the same type as the first argument, but by parametricity we can deduce that the result will be the first argument. Likewise for the second argument. A simple way to understand it, is by instantiating both type variables in (2) with the singleton type of the (value result of the) argument.

then, as we explain later on, our system infers that `id` has type $\forall \alpha. \alpha \rightarrow \alpha$, viz., that `id` is indeed the polymorphic identity function.

This is clearly better than the current state of the art. Still, it does not seem too hard a feat to deduce that if we are testing whether `x` is a truthy value, then when the test succeeds we can assume that `x` is of type `Truthy`. To show the more advanced capabilities of our system let us have a look at how ECMAScript specifies the semantics of JavaScript logical operators, as defined in the 2021 version of the specification [Ecma 2021, Section 13.13.1]. Since in JavaScript there are no union or intersection types, then the falsy and truthy values are defined via an (abstract) function `ToBoolean` which simply checks whether its argument is one of the 8 falsy values and returns `false`, otherwise it returns `true` (see its definition in row 1 of Table 1 in Section 5). In our system, `ToBoolean` has type $(\text{Truthy} \rightarrow \text{true}) \wedge (\text{Falsy} \rightarrow \text{false})$. All logical operators are then defined by ECMAScript in terms of this function: this has the advantage that any change to the specification of falsy (e.g., the addition of a new falsy value, like the addition of the built-in `bigint` type and its constant `0n` in ES2020) requires only the modification of this function, and is automatically propagated to all operators. So the actual definition of `10r` for ECMAScript is the following one:

```

7  function 10r (x, y) {
8      if (ToBoolean(x)) { return x; } else { return y; }
9  }

```

If we feed this function to our system, then it infers for it the type in (2), that is, the same type it already deduced for the simpler version of `10r` defined in lines 1-3. But here the deduction needed to perform type narrowing is more challenging, since the system must deduce from the type $(\text{Truthy} \rightarrow \text{true}) \wedge (\text{Falsy} \rightarrow \text{false})$ of `ToBoolean` that when the *application* in line 8 returns a truthy value, then the *argument* of `ToBoolean` is of type `Truthy`, and it is of type `Falsy` otherwise. More generally, we need a system which, when a test is performed on an arbitrarily complex application, can narrow the type of all the variables occurring in the application by exploiting the information provided by the overloaded behavior of the functions therein. Achieving such a degree of precision is a hard feat but, we argue, it is necessary if we want to reconstruct types for dynamic languages, that is, if we want to type their programs as they are, without requiring the addition of any type annotations. Indeed, the core operators of these languages (e.g., JavaScript's "`||`", "`&&`", "`typeof`", ...) are characterized by an "overloaded" behavior, which is then passed over to the functions that use them. So for instance a simple use of JavaScript logical or "`||`" such as in `(x => x || 42)` results in a function whose precise type, as reconstructed by our system, is $(\text{Falsy} \rightarrow 42) \wedge (\text{Truthy} \wedge \alpha \rightarrow \text{Truthy} \wedge \alpha)$. JavaScript functions also routinely perform dynamic checks against constants (notably `null` and `undefined`), which our system also handles as part of its more general approach to type narrowing of arbitrary expressions.

1.1 Outline

Type System (Section 2). So, how can we achieve all this? Conceptually, it is quite simple: we just merge together three of the most expressive type systems studied in the literature, namely the Hindley-Milner (HM) polymorphic types [Hindley 1969; Milner 1978], intersection types [Coppo et al. 1981], and union types [Barbanera et al. 1995; MacQueen et al. 1986]. We achieve it simply by putting together in a controlled way the deduction rules characteristic of each of these systems (see Figure 2 in Section 2) and proving that the resulting system is sound (cf., Theorem 2.2).

More precisely, the type system we describe in Section 2 is pretty straightforward. Its core is a classic HM system with first order polymorphism: a program is a list of `let`-bindings that define polymorphic functions; these are typed by inferring a type for the expressions that define them, this type is then generalized, yielding a prenex polymorphic type for the function. As usual,

the deduction of the type of each of these expressions is performed in a type environment that records the generic types for the previously-defined polymorphic functions, and the type system can instantiate these types differently for each use of the polymorphic functions in the expression. The novelty of our system is that when deducing the types of the expressions that define the polymorphic functions, the type system can use not only instantiations of polymorphic types (rule [INST] in Figure 2), but also intersection and union types. More precisely, to type these expressions the type system can decide to use the classic rules of intersection introduction (rule $[\wedge]$) and union elimination (rule $[\vee]$) given in Figure 2. For instance, the intersection introduction rule is used by the system to deduce that since the function `10r` (either versions) has both type $((\alpha \wedge \text{Truthy}, \text{Any}) \rightarrow \alpha \wedge \text{Truthy})$ and type $((\text{Falsy}, \beta) \rightarrow \beta)$, then it has their intersection, too; this intersection type is then generalized (when `10r` is defined at top-level) yielding the polymorphic type in (2). The union elimination rule is essentially used to fine-grainedly type branching expressions and tests involving applications of overloaded functions: for instance, to deduce that the function `id` in lines 4–6 has type $\alpha \rightarrow \alpha$, the system can assume that `x` has type α and separately infer the type of the body for `x` : $(\alpha \wedge \text{Truthy})$ and for `x` : $(\alpha \wedge \neg \text{Truthy})$; since the first deduction yields $(\alpha \wedge \text{Truthy})$ and the second yields $(\alpha \wedge \neg \text{Truthy})$, then the system deduces that under the hypothesis `x` : α , the body has the union of these two types, that is α . Furthermore, as observed by Castagna et al. [2022b], the combination of the union elimination with the rules of type-cases given in Figure 2 constitutes the essence of narrowing and occurrence typing.

The declarative type system given in Section 2 is all well and good, but how can we define an algorithm that infers whether a given expression can be typed in this system? Rules such as union elimination and intersection introduction are easy to understand, but they do not easily lend themselves to an implementation. In order to arrive to an effective implementation of the type system specified in Section 2 we proceed in two steps: (i) the definition of an algorithmic system and (ii) the definition of a reconstruction algorithm.

Algorithmic System (Section 3). The first step towards an effective implementation of our type system is taken in Section 3 where we define an algorithmic system that is sound and complete with respect to the system of Section 2. The system is algorithmic since it is composed only by syntax-directed and analytic rules⁴ and, as such, is immediately implementable. It is sound and complete since an expression is typable in it if and only if it is typable in the system of Section 2. To obtain this results the system is defined on pairs formed by an MSC-form (Maximal Sharing Canonical form) and an annotation tree. MSC-forms are A-normal forms [Sabry and Felleisen 1992] on steroids: they are lists of bindings associating variables to expressions in which every proper subexpression is a variable. Their characteristic is that they encode expressions and preserve typability in the sense that every expression is typable if and only if its unique MSC-form is typable. MSC-forms were introduced by Castagna et al. [2022b] to drastically reduce the range of possible applications of the union elimination rule; here we improve their definition to deal with our polymorphic setting and use them for exactly the same reason as in [Castagna et al. 2022b]. Annotation trees encode canonical derivations of the system of Section 2 for the MSC-form they are paired with. They are a generalization of type annotations inserted in the code. Instead of annotating directly an MSC-form with type-annotations we used a separate annotation tree because of the union elimination rule which types several times the same expression under different type environments; this would, thus, require different annotations for the same subexpressions, each annotation depending on the typing context: this naturally yields to tree-shaped annotations in which each branching corresponds either to the different deductions performed by a union elimination rule or to the different deductions

⁴A rule is *analytic* (as opposed to *synthetic*) when the input (i.e., Γ and e) of the judgment at the conclusion is sufficient to determine the inputs of the judgments at the premises (cf. [Martin-Löf 1994; Types 2019]).

performed by an intersection introduction rule. The soundness and completeness properties of the algorithmic systems are thus stated in terms of MSC-forms and annotation trees. They essentially state that an expression e has type t in the declarative system of Section 2 if and only if there exists a tree annotation for the (unique) MSC-form of e that is typable in the algorithmic system with (a subtype of) t : see Theorem 3.4.

Reconstruction Algorithm (Section 4). The second of the two steps to achieve an effective implementation for the type system of Section 2 is to define a reconstruction algorithm for the previous algorithmic system, which we do in Section 4. The statements of the soundness and completeness properties of the algorithmic system clearly suggest what this algorithm is expected to do: given an expression that defines a polymorphic function, the algorithm must transform it into its unique MSC-form and then try to reconstruct an annotation tree for it so that the pair MSC-form and annotation tree is typable in the algorithmic system of Section 3.

The reconstruction is performed by a system of deduction rules that incrementally refines an annotation tree (initially composed of a single node “infer”) while exploring the list of bindings of the MSC-form of the expression to type. It mixes two independent mechanisms: one that infers the domain(s) of λ -terms, and the other that performs type narrowing when a typecase is encountered.

The first mechanism is inspired by the algorithm \mathcal{W} by Damas and Milner [1982]: whenever the application of a destructor (e.g., a function application) is encountered, an algorithm finds a substitution (if any) that makes this application well-typed. In the context of a HM type system, the algorithm at issue needs to solve a *unification problem* (i.e., whether for two given types s and t there exists a substitution σ such as $s\sigma = t\sigma$) which, if solvable, has a principal solution given by a single substitution [Robinson 1965]. In our system, which is based on subtyping, the algorithm at issue needs to solve a *tallying problem* (i.e., whether for two given types s and t there exists a substitution σ such as $s\sigma \leq t\sigma$) which, if solvable, has a principal solution given by a *finite set* of substitutions [Castagna et al. 2015]. When multiple substitutions are found, they are all considered and explored in different branches by adding an intersection branching node in the current annotation tree.

The second mechanism gets inspiration from Castagna et al. [2022b] and refines decompositions made by the union-elimination rule in order to narrow the types of variables in the branches of a typecase expression. When the system encounters a typecase that tests whether some expression e has type t , then the type s of the variable bound to e (recall that an MSC-form is a list of bindings) is split into $s \wedge t$ and $s \wedge \neg t$, and these splits are in turn propagated recursively in order to generate new splits for the types of the variables associated with the subexpressions composing e . For instance, when the algorithm encounters the test “**if** (ToBoolean(x)) . . .” at line 8, it splits the type of (the variable bound to) ToBoolean(x) in two, by intersecting it with **true** and \neg **true**, and this split in turn generates a new split **Truthy** and **Falsy** for the type of the variable x .

The reconstruction algorithm we present in Section 4 is sound: if it returns an annotation tree for an MSC-form, then the pair is typable in the algorithmic system, whose soundness implies that the expression at the origin of the MSC-form is typable in the system of Section 2. At this point, however, it should be pretty obvious that such a reconstruction algorithm cannot be complete. Our system merges three well known systems: first-order parametric polymorphism, intersection types, union elimination. Now, even if parametric polymorphism is decidable, in our system we can encode (and type, via intersection types) polymorphic fixed-point combinators, yielding a system with polymorphic recursion whose inference has been long known to be undecidable [Henglein 1993; Kfoury et al. 1993]. Worse, our system includes union elimination, which is one of the most problematic rules from an algorithmic viewpoint, not only because it is neither syntax directed nor analytic, but also because determining an inversion (a.k.a., generation) lemma for this rule is

considered by experts the most important open problem in the research on union and intersection types [Dezani 2020], and an inversion lemma is somehow the first step to define a type-inference algorithm, since it tells us when and how to apply the rule. We discuss in detail the reasons and implications of incompleteness in Section 4.4.

Despite being incomplete, our reconstruction algorithm is powerful enough to handle both complicated typing use-cases and common programming patterns of dynamic languages. For instance, for the Z fixed-point combinator for strict languages $Z = \lambda f.(\lambda x.f(\lambda v.xxv)) (\lambda x.f(\lambda v.xxv))$ our algorithm reconstructs the type $\forall \alpha, \beta, \gamma. ((\alpha \rightarrow \beta) \rightarrow ((\alpha \rightarrow \beta) \wedge \gamma)) \rightarrow ((\alpha \rightarrow \beta) \wedge \gamma)$ (i.e., in bounded polymorphic notation $\forall(\alpha)(\beta)(\gamma \leq \alpha \rightarrow \beta).((\alpha \rightarrow \beta) \rightarrow \gamma) \rightarrow \gamma$, cf. Footnote 2). The combinator can then be used as is, to define and infer the type of classic polymorphic functions such as `map`, `fold`, `concat`, `reverse`, etc., often yielding types more precise than in HM: for instance if we use $[\alpha^*]$ to denote the type of the lists whose elements have type α , then the type inferred for (a curried version of) `fold_r` is $\forall \alpha, \beta, \gamma. ((\alpha \rightarrow \beta \rightarrow \beta) \rightarrow \beta \rightarrow [\alpha^*] \rightarrow \beta) \wedge (\text{Any} \rightarrow \gamma \rightarrow [] \rightarrow \gamma)$ where the second type in the intersection states that if the third argument is an empty list, then the result will be the second argument, whatever the type of the first argument is. Finally, we designed our algorithm so that it can take into account explicit type annotations to help it in the inference process. As an example, our algorithm can check that the classic `filter` function has type $\forall \alpha, \beta, \gamma. ((\alpha \wedge \beta \rightarrow \text{Bool}) \wedge (\alpha \wedge \neg \beta \rightarrow \text{False})) \rightarrow [\alpha^*] \rightarrow [(\alpha \wedge \beta)^*]$, stating that if we pass to `filter` a predicate that returns false for the elements of α that are not in β , then filtering a list of α 's will return only elements also in β .

Sections 2, 3, and 4 outlined above constitute the core of our contribution. Section 5 presents our implementation. In Section 6 we discuss related work and Section 7 concludes our presentation. For space reasons we omitted in the main text some rules of the algorithmic and reconstruction systems, as well as all proofs: they are all given in the appendix, available on line as supplemental material [Castagna et al. 2024].

1.2 Discussion, Contributions, and Limitations

Intersections vs. Hindley-Milner. It is a truth universally acknowledged that intersection type systems are more powerful than HM systems: for that, one does not even need full intersections, since Rank 2 intersections suffice. Rank 2 intersection types are types that may contain intersections only to the left of a single arrow and the system of Rank 2 intersection types is able to type all ML programs (i.e., all program typable by HM), has principal typings, decidable type inference, and the complexity of type inference is of the same order as in ML [Leivant 1983].

However, intersection type systems are not compositional, and this hinders their use in a modular setting. A program that uses the polymorphic identity function to apply it to, say, an integer and a Boolean, type checks since we can infer that the polymorphic identity function has type $(\text{Int} \rightarrow \text{Int}) \wedge (\text{Bool} \rightarrow \text{Bool})$. But if we want to *export* this polymorphic identity function to use it in other unforeseen contexts, then we need for it a type that covers all its admissible usages, without the need of retype-checking the function every time it is applied to an argument of a new type. In other words, in a modular usage, parametric polymorphism has an edge over intersection/ad-hoc polymorphism despite being less powerful, since a type such as $\forall \alpha. \alpha \rightarrow \alpha$ synthesizes the infinitely many combinations of intersection types that can be deduced for the identity function; however in a local setting, everything that does not need to be exported can be finer-grainedly typed by intersection types. This division of roles and responsibilities is at the core of our approach. As we show in the next section, programs are lists of bindings from variables to expressions. These expressions are typed in a type environment (generated by the preceding bindings) which binds variables to polymorphic types. These expressions are typed by using instantiation, intersection

introduction, and union elimination, but *not* generalization. Generalization is performed only at top level, that is at the level of programs and reserved to expressions to be used in other contexts.

Parametricity vs. type cases. A parametric polymorphic function (a.k.a., a generic function) is a function that behaves uniformly for all possible types of its arguments, that is, whose behavior does not depend on the type of its arguments. A common way to characterize a generic function is that it is a function that cannot inspect the parametric parts of its input, that is, those parts that are typed by a type variable: these parts can only be either returned as they are, or discarded, or passed to another generic function. Our approach suggests refining this characterization by shifting the attention from inputs as a whole to some particular values among all the possible inputs. This can be seen by comparing the following two function definitions:

$$\lambda x. (x \in \text{Int}) ? x : x \qquad \lambda x. (x \in \text{Int}) ? x + 1 : x$$

Both functions test whether their input is an integer. The function on the right-hand side returns the successor of the argument when this is true and the argument itself otherwise; the one on the left-hand side returns its argument in both cases, that is, it is the identity function.

Our system deduces for the function on the left the type $\alpha \rightarrow \alpha$.⁵ For the function on the right it returns the type $(\text{Int} \rightarrow \text{Int}) \wedge (\alpha \setminus \text{Int} \rightarrow \alpha \setminus \text{Int})$, where $s \setminus t$ denotes the set-theoretic difference of the two types, that is, $s \wedge \neg t$. These two types suggest how we can refine the intuitive characterization of parametricity. A generic function can inspect the parametric part of its input (as the function on the left-hand side shows) and its output can depend on this inspection (as the function on the right-hand side shows), but the parts of its output that are typed by a type variable—i.e., the “parametric” parts—cannot depend on it. We can speak of “partial” parametricity, and say that a function is parametric “only” for the inputs (or parts thereof) that are either returned unchanged or discarded: the type variables in its type describe such inputs. For instance, the domain of both the functions above is Any : they both can be applied to any argument. But the first function is parametric for all possible inputs, since the result of the inspection is not used to produce any particular output (it has type $\forall \alpha. \alpha \rightarrow \alpha$), while the second function is parametric only for the values of its domain that are not in Int , since it uses the result of the inspection to generate the result for the integer inputs (by subsumption, the second function has type $\forall \alpha. \alpha \rightarrow \text{Int} \vee (\alpha \setminus \text{Int})$: parametricity holds only for the arguments not in Int).

Contributions. The general contribution of our work is twofold. First, it proposes a way to mix parametric and intersection/ad-hoc polymorphism which, in hindsight, is natural: parametric polymorphism for everything defined at top-level and that can thus be used in other contexts (modularity); intersection polymorphism for everything that remains local (for which we can thus use more precise non-modular typing). Second, it proposes an effective way to implement this type discipline by defining a reconstruction algorithm; with respect to that, a fundamental role is played by the analysis of the (type-)tests performed by the expressions, since they drive the way in which types are split: externally, to split the domain of functions yielding intersection of arrows (intersection introduction); internally, to split the type of tested expressions, yielding a precise typing of branching (union elimination). In doing so, it provides the first system that reconstructs types that oncombine parametric and *ad hoc* polymorphism.

The technical contributions of the work can be summarized as follows:

- (1) We define a type system that combines parametric polymorphism with union and intersection types for a functional calculus with type-cases and prove its soundness.
- (2) We define an algorithmic system that we prove sound and complete with respect to the previous system.

⁵Precisely, it deduces for it the type $(\alpha \setminus \text{Int} \rightarrow \alpha \setminus \text{Int}) \wedge (\beta \setminus \text{Int} \rightarrow \beta \setminus \text{Int})$. Instantiating β to α yields a subtype of $\alpha \rightarrow \alpha$.

- (3) We define an algorithm to reconstruct the type annotations of the previous algorithmic system and prove it sound and terminating.

The reconstruction algorithm is fully implemented. A prototype which also implements optional type annotations and pattern matching (presented in Appendix A) is available on-line at <https://www.cduce.org/dynlang>, and whose sources are on Zenodo: [Castagna et al. 2023b].

Limitations. The system we present here has some limitations. Foremost, the reconstruction algorithm of Section 4 uses backtracking, and at each of its passes it may try to type the same piece of code several times. Backtracking is inherent to our algorithm, since it proceeds by successively refining in different passes, the annotation tree of an MSC-form. The checking of a same piece of code several times at each pass is inherent to the use of unions and intersections: the union-elimination rule repeatedly type-checks the same expression, using different type hypotheses for a given sub-expression; the intersection-introduction rule verifies that an expression has all the types of an intersection, by checking each of them separately. Both features are very penalizing in terms of performance, and any naive implementation of the reconstruction described in Section 4 would yield type-inference times that grow exponentially with the size of the program. Clearly, this is an issue that must be addressed if we want to apply our system to real-world dynamic languages, and further work is needed to frame and/or constrain the current system so that its performance becomes acceptable. Fortunately, the room for improvement is significant: our prototype is an unoptimized proof of concept whose implementation was defined to faithfully simulate the reconstruction inference rules, rather than to obtain an efficient execution; but the simple addition of textbook memoization techniques improved its performance by an order of magnitude (cf. Section 5).

A second limitation of our system is that it is not sound in the presence of side-effects. The algorithm transforms an initial expression into its Maximal Sharing Canonical form, which is a list of bindings, one for each sub-expression of the initial expression. As we explain in Section 3.1.2, these forms are called “maximal sharing” since all equivalent sub-expressions (in the sense stated by Definition 3.1) of the initial expression must be bound by the same variable, so that any refinement of the type of one sub-expression (e.g., as a consequence of a type-case) is passed-through to all equivalent sub-expressions. However, this is sound only if all evaluations of equivalent sub-expressions return results that have the same types. While this is true for pure expressions, this can be invalidated by the presence of side-effects. In Section 7 we suggest some research directions on how to modify the equivalence relation of Definition 3.1 to make our system work in the presence of side-effects. Nevertheless, the work presented here is closer to be adapted/adaptable to pure functional languages such as Erlang and Elixir, than to languages such as JavaScript or Python.

Finally, it may be worth pointing out that our approach works only for strict languages, since it uses a semantic subtyping relation that is unsound for call-by-name evaluation strategies [Petrucci et al. 2018].

2 SOURCE LANGUAGE AND TYPE SYSTEM

2.1 Syntax and Semantics

Our core language is fully defined in Figure 1. Expressions are an untyped λ -calculus with constants c , pairs (e, e) , pair projections $\pi_i e$, and type-cases. A type-case $(e_0 \in \tau) ? e_1 : e_2$ is a dynamic type test that first evaluates e_0 and, then, if e_0 reduces to a value v , evaluates e_1 if v has type τ or e_2 otherwise. Type-cases cannot test arbitrary types but just ground types (i.e., types without type variables occurring in them) of the form τ where the only arrow type that can occur in them is $\mathbb{0} \rightarrow \mathbb{1}$, the type of all functions. This means that type-cases can distinguish functions from other values but they cannot distinguish, say, functions that have type $\text{Int} \rightarrow \text{Int}$ from those that do not.

Syntax

Test Type $\tau ::= b \mid \mathbb{0} \rightarrow \mathbb{1} \mid \tau \times \tau \mid \tau \vee \tau \mid \neg \tau \mid \mathbb{0}$
Expression $e ::= c \mid x \mid \lambda x. e \mid ee$
 $\quad \mid (e, e) \mid \pi_i e \mid (e \in \tau) ? e : e$
Value $v ::= c \mid \lambda x. e \mid (v, v)$
Program $p ::= \text{let } x = e ; p \mid e$

Reduction rules

$(\lambda x. e)v \rightsquigarrow e\{v/x\}$
 $\pi_1(v_1, v_2) \rightsquigarrow v_1$
 $\pi_2(v_1, v_2) \rightsquigarrow v_2$
 $(v \in \tau) ? e_1 : e_2 \rightsquigarrow e_1 \quad \text{if } v \in \tau$
 $(v \in \tau) ? e_1 : e_2 \rightsquigarrow e_2 \quad \text{if } v \in \neg \tau$
 $\text{let } x = v ; p \rightsquigarrow_{pr} p\{v/x\}$

Dynamic type test

$$v \in \tau \Leftrightarrow \text{typeof}(v) \leq \tau, \text{ where } \begin{cases} \text{typeof}(c) & = b_c \\ \text{typeof}((v_1, v_2)) & = \text{typeof}(v_1) \times \text{typeof}(v_2) \\ \text{typeof}(\lambda x. e) & = \mathbb{0} \rightarrow \mathbb{1} \end{cases}$$
Evaluation Contexts

$E ::= [] \mid Ee \mid vE \mid (E, e) \mid (v, E) \mid \pi_i E \mid (E \in \tau) ? e : e$
 $P ::= [] \mid \text{let } x = [] ; p$

$$\frac{e \rightsquigarrow e'}{E[e] \rightsquigarrow E[e']} \quad \frac{e \rightsquigarrow e'}{P[e] \rightsquigarrow_{pr} P[e']}$$

Fig. 1. Syntax and semantics of the source language

Programs are sequences of top-level definitions, ending with an expression that can be seen as the main entry. This notion of program is useful to capture the modularity of our type system. Indeed, top-level definitions are typed sequentially: the type we obtain for a top-level definition is considered definitive and will not be challenged by a later definition.

The reduction semantics for expressions is the one of call-by-value pure λ -calculus with products and with a type-case expression, together with the context rules that implement a leftmost outermost reduction strategy. We use the standard substitution operation $e\{e'/x\}$ that denotes the capture avoiding substitution of e' for x in e , whose definition we recall in Appendix B. The relation $v \in \tau$ determines whether a *value* is of a given type or not and holds true if and only if $\text{typeof}(v) \leq \tau$, where \leq is the subtyping relation defined by Castagna and Xu [2011] (we recall its definition in Appendix C). Note that $\text{typeof}(v)$ maps every λ -abstraction to $\mathbb{0} \rightarrow \mathbb{1}$ and, thus, dynamic type tests do not depend on static type inference. This approximation is allowed by the restriction on arrow types in typecases. Finally, the reduction semantics for programs sequentially reduces top-level definitions, together with a context rule that allows reducing the expression of the first definition.

2.2 Types

Types are those by Castagna and Xu [2011] who add type variables to the semantic subtyping framework of Frisch et al. [2002, 2008].

DEFINITION 2.1 (TYPES). *The set of types **Types** is formed by the terms t coinductively produced by the grammar:*

$$\mathbf{Types} \quad t, s ::= b \mid \alpha \mid \alpha \mid t \rightarrow t \mid t \times t \mid t \vee t \mid \neg t \mid \mathbb{0}$$

and that satisfy the following conditions: (i) every term has a finite number of different sub-terms (regularity) and (ii) every infinite branch of a term contains an infinite number of occurrences of the arrow or product type constructors (contractivity).

We use the abbreviations $t_1 \wedge t_2 \stackrel{\text{def}}{=} \neg(\neg t_1 \vee \neg t_2)$, $t_1 \vee t_2 \stackrel{\text{def}}{=} t_1 \wedge (\neg t_2)$, and $\mathbb{1} \stackrel{\text{def}}{=} \neg \mathbb{0}$. Basic types (e.g., Int, Bool) are ranged over by b , $\mathbb{0}$ and $\mathbb{1}$ respectively denote the empty (that types no value) and top (that types all values) types. Coinduction accounts for recursive types and the condition on infinite branches bars out ill-formed types such as $t = t \vee t$ (which does not carry any information about the set denoted by the type) or $t = \neg t$ (which cannot represent any set).

For what concerns type variables, we choose *not* to use type-schemes but rather distinguish two kinds of type variables. *Polymorphic type variables* ranged over by α , are type variables that have been generalized and can therefore be instantiated. In a more traditional presentation, such variables are bound by the \forall of a type-scheme; the set of polymorphic variables is \mathcal{V}_P . *Monomorphic type variables*, ranged over by α (with bold font), are variables that are not generalized and therefore cannot be instantiated; the set of monomorphic variables is \mathcal{V}_M . Types that only contain monomorphic variables are dubbed monomorphic types⁶:

$$\mathbf{Monomorphic\ types} \quad \mathbf{u, v} ::= b \mid \alpha \mid \mathbf{u} \rightarrow \mathbf{u} \mid \mathbf{u} \times \mathbf{u} \mid \mathbf{u} \vee \mathbf{u} \mid \neg \mathbf{u} \mid \mathbb{0}$$

Our choice of using two disjoint sets for polymorphic and monomorphic type variables, instead of the classical approach of using type schemes $\forall \alpha_1 \dots \alpha_n. t$, is justified by two reasons. First, type schemes are expected to be equivalent modulo α -renaming. In our case however, we do not want polymorphic type variables to be freely renamed because of the use, in the algorithmic type system of Section 3, of external annotations containing explicit substitutions over some polymorphic type variables of the context. Secondly, introducing type schemes would require redefining many of the usual set-theoretic type-related definitions, such as the subtyping relation \leq , and the type operators for application \circ and projections π_i . Instead, we obtain a more streamlined theory by making subtyping and these operators ignore whether a variable is polymorphic or monomorphic in the current context, and by explicitly performing instantiations in the type system when required.

The subtyping relation for these types, noted \leq , is the one defined by Castagna and Xu [2011], to which the reader may refer for the formal definition (cf. Appendix C). For this presentation, it suffices to consider that ground types (i.e., types with no variables) are interpreted as sets of *values* that have that type, and that subtyping is set containment (i.e., a type s is a subtype of a type t if and only if t contains all the values of type s). In particular, $s \rightarrow t$ contains all λ -abstractions that when applied to a value of type s , if their computation terminates, then they return a result of type t (e.g., $\mathbb{0} \rightarrow \mathbb{1}$ is the set of all functions and $\mathbb{1} \rightarrow \mathbb{0}$ is the set of functions that diverge on every argument). Type connectives (i.e., union, intersection, negation) are interpreted as the corresponding set-theoretic operators. For what concerns non-ground types (i.e., types with variables occurring in them) all the reader needs to know for this work is that the subtyping relation of Castagna and Xu [2011] is preserved by type-substitutions. Namely, if $s \leq t$, then $s\sigma \leq t\sigma$ for every type-substitution σ . We use \simeq to denote the symmetric closure of \leq , thus $s \simeq t$ (read, s is equivalent to t) means that s and t denote the same set of values and, as such, they are semantically the same type.

2.3 Type System

Our type system is given in full in Figure 2. The typing rules for expressions are, to some extent, the usual ones. Constants and variables are typed by the corresponding axioms [CONST] and [Ax]. The arrow and product constructors have introduction and elimination rules. Notably, in the case of rule $[\rightarrow I]$ the type of the argument is monomorphic. The rules for intersection ($[\wedge]$) and subtyping ($[\leq]$) are the classical ones, and so is the rule for instantiation ([INST]) where σ denotes a substitution from polymorphic variables to types. The type-case construction is handled by three rules: $[\mathbb{0}]$; $[\in_1]$; $[\in_2]$. Rule $[\mathbb{0}]$ handles the case where the tested expression is known to have the empty type. The other two are symmetric and handle the case when the tested expression is known

⁶The term *polytypes* and *monotypes* can be found (albeit inconsistently) in the literature: in particular, Milner [1978] uses the latter to denote types with no type variables and the former when he wishes to imply that a type may, or does, contain a variable. We avoided using them to prevent any confusion with our *monomorphic types*. While our *types* are indeed polytypes, our *monomorphic types* are not monotypes: monotypes do not have type variables while monomorphic types may have type variables, though only monomorphic ones. So we used instead *types* (which may have type variables), *ground types* (which cannot have any type variable), and *monomorphic types* (which may have monomorphic type variables, only).

$$\begin{array}{c}
\text{[CONST]} \frac{}{\Gamma \vdash c : b_c} \qquad \text{[Ax]} \frac{}{\Gamma \vdash x : \Gamma(x)} \\
\text{[}\rightarrow\text{I]} \frac{\Gamma, x : \mathbf{u} \vdash e : t}{\Gamma \vdash \lambda x. e : \mathbf{u} \rightarrow t} \qquad \text{[}\rightarrow\text{E]} \frac{\Gamma \vdash e_1 : t_1 \rightarrow t_2 \quad \Gamma \vdash e_2 : t_1}{\Gamma \vdash e_1 e_2 : t_2} \\
\text{[}\times\text{I]} \frac{\Gamma \vdash e_1 : t_1 \quad \Gamma \vdash e_2 : t_2}{\Gamma \vdash (e_1, e_2) : t_1 \times t_2} \qquad \text{[}\times\text{E}_1\text{]} \frac{\Gamma \vdash e : t_1 \times t_2}{\Gamma \vdash \pi_1 e : t_1} \qquad \text{[}\times\text{E}_2\text{]} \frac{\Gamma \vdash e : t_1 \times t_2}{\Gamma \vdash \pi_2 e : t_2} \\
\text{[0]} \frac{\Gamma \vdash e : \emptyset}{\Gamma \vdash (e \in \tau) ? e_1 : e_2 : \emptyset} \qquad \text{[}\epsilon_1\text{]} \frac{\Gamma \vdash e : \tau \quad \Gamma \vdash e_1 : t_1}{\Gamma \vdash (e \in \tau) ? e_1 : e_2 : t_1} \qquad \text{[}\epsilon_2\text{]} \frac{\Gamma \vdash e : \neg \tau \quad \Gamma \vdash e_2 : t_2}{\Gamma \vdash (e \in \tau) ? e_1 : e_2 : t_2} \\
\text{[}\vee\text{]} \frac{\Gamma \vdash e' : s \quad \Gamma, x : s \wedge \mathbf{u} \vdash e : t \quad \Gamma, x : s \wedge \neg \mathbf{u} \vdash e : t}{\Gamma \vdash e\{e'/x\} : t} \qquad \text{[}\wedge\text{]} \frac{\Gamma \vdash e : t_1 \quad \Gamma \vdash e : t_2}{\Gamma \vdash e : t_1 \wedge t_2} \\
\text{[INST]} \frac{\Gamma \vdash e : t}{\Gamma \vdash e : t\sigma} \qquad \text{[}\leq\text{]} \frac{\Gamma \vdash e : t}{\Gamma \vdash e : t'} \quad t \leq t'
\end{array}$$

Fig. 2. Declarative Type System

to have either the type τ or its negation, in which case the corresponding branch is typed. These rules work together with Rule [V], which we describe in detail next.

At first sight, the formulation of rule [V] seems odd, since the \vee connector does not appear in it. To understand it, consider the classic union elimination rule by [MacQueen et al. \[1986\]](#):

$$\text{[VE]} \frac{\Gamma \vdash e' : s_1 \vee s_2 \quad \Gamma, x : s_1 \vdash e : t \quad \Gamma, x : s_2 \vdash e : t}{\Gamma \vdash e\{e'/x\} : t}$$

Rule [VE] types an expression that contains occurrences of an expression e' that has a union type $s_1 \vee s_2$; the rule substitutes in this expression *some* occurrences of e' by the variable x yielding an expression e , and then types e first under the hypothesis that x has type s_1 and then under the hypothesis that x has type s_2 . If both succeed, then the common type is returned for the expression at issue. As shown by [Castagna et al. \[2022b\]](#), this rule, together with the rules for type-cases, allows the system to perform occurrence typing. For instance, consider the expression $(fy \in \text{Int}) ? (fy) + 1 : \mathbf{false}$, in the context where f has type $\text{Any} \rightarrow \text{Any}$ and y is of type Any . This expression can be typed thanks to the rule [VE], by considering the sub-expression fy . This sub-expression has type Any , which can be seen as the union type $\text{Any} \simeq \text{Int} \vee \neg \text{Int}$. We can then replace x for fy and type, using $[\epsilon_1]$, the expression $(x \in \text{Int}) ? x + 1 : \mathbf{false}$, with $x : \text{Int}$. This yields a type Int (rule $[\epsilon_1]$ ignores the second branch) and by subtyping, the expression has type $\text{Int} \vee \mathbf{False}$. Likewise for the choice $x : \neg \text{Int}$, using rule $[\epsilon_2]$ the second branch has type \mathbf{False} and therefore $\text{Int} \vee \mathbf{False}$ (again via subtyping). The whole expression has thus the desired type $\text{Int} \vee \mathbf{False}$.

A key element is that rule [VE] guessed how to split the type Any of fy into $\text{Int} \vee \neg \text{Int}$. In a non-polymorphic setting, this is perfectly fine. But in a type-system featuring polymorphism, particular care must be taken when introducing (fresh) type variables. As it is stated, MacQueen et al.'s [VE] rule could choose to split, say, $\mathbb{1}$ into a union $\alpha \vee \neg \alpha$, with α a polymorphic type variable. If so, then the rule becomes unsound. As a matter of fact, the premises of the [VE] behave as in rule $[\rightarrow\text{I}]$, in that they introduce in the typing environment a fresh type whose variables must *not* be instantiated. In our example, however, in one premise, the rule introduces $x : \alpha$ in the typing

environment which can, for instance, be instantiated by the [Inst] rule. In the second premise, it introduces $x : \neg\alpha$ which can also be instantiated in a *completely different way*. In other words, the correlation between the two occurrences of the same variable α is lost, which amounts to commuting the (implicit) universal quantification with the \vee type connective, yielding a non-prenex polymorphic type $(\forall\alpha.\alpha) \vee (\forall\alpha.\neg\alpha)$. To avoid this unsound situation, we need to ensure that when a type is split between two components of a union, no polymorphic variable is introduced. This is achieved by rule [V] which requires the type s of e' to be split as $s \equiv (s \wedge \mathbf{u}) \vee (s \wedge \neg\mathbf{u})$ (here is our hidden union).

The top-level definitions of a program are typed sequentially by two specific rules:

$$\begin{array}{c} \text{[TOPLEVEL-EXPR]} \quad \frac{\Gamma \vdash e : t}{\Gamma \vdash_{Pr} e : t\phi} \phi\#\Gamma \quad \text{[TOPLEVEL-LET]} \quad \frac{\Gamma \vdash_{Pr} e : t \quad \Gamma, x : t \vdash_{Pr} p : t'}{\Gamma \vdash_{Pr} \text{let } x = e ; p : t'} \end{array}$$

where ϕ denotes a generalization, that is a substitution transforming monomorphic variables into polymorphic ones and where $\phi\#\Gamma \stackrel{\text{def}}{\iff} \text{dom}(\phi) \cap \text{vars}(\Gamma) = \emptyset$.

After typing an expression used for a definition, its type is generalized (Rule [TOPLEVEL-EXPR]) before being added in the environment (Rule [TOPLEVEL-LET]). Note that this is the only place where generalization takes place: no rule in the type system for expressions (Figure 2) allows the generalization of a type variable. As explained at the beginning of Section 1.2, this is not a limitation, since intersection types are more powerful than HM polymorphism, and top-level generalization is of practical importance since it is necessary to the modularity of type-checking. Nevertheless, the core of our inference system is given only by the rules in Figure 2 for expressions: the above ‘‘TOPLEVEL’’ rules are only useful to inhabit variables of the typing environments used in the rules for expressions, and this makes it possible to close the expressions being typed. For instance, if a typing derivation for an expression e is deduced, say, under the hypothesis $x : \alpha \rightarrow \alpha$ (with α polymorphic), then it is possible to obtain a closed program by inhabiting x by a definition like $\text{let } x = \lambda y.y ; \dots ; e$. This is the reason why, henceforth, we mainly focus on the typing of expressions.

The type system is sound (all proofs for this work are given in Appendix I):

THEOREM 2.2 (SOUNDESS). *If $\emptyset \vdash_{Pr} p : \tau$, then either p diverges or $p \rightsquigarrow_{Pr} v$ with $v \in \tau$.*

3 ALGORITHMIC SYSTEM

As discussed in the introduction, the declarative type system is not syntax directed and some rules are not analytic. In order to make it algorithmic, we first introduce in Section 3.1 a *canonical form* for expressions that adds syntactic constructions (*bindings*) to indicate when to apply the union elimination rule and on which sub-expression. Then, in Section 3.2, we define a fully algorithmic type system that takes a *canonical form* together with an *annotation* and produces a type.

3.1 MSC Forms

3.1.1 Canonical Forms. The [V] rule is not syntax directed since it can be applied on any expression and can split the type of any of its subexpressions. If we want an algorithmic type system, we need a syntactic way to determine when to apply this rule, and which subexpression to split. In order to achieve this, we represent our terms with a syntax called *Maximal Sharing Canonical Form* (MSC Form) introduced by Castagna et al. [2022b]. Let us start by defining the *canonical forms*, which are expressions produced by the following grammar:

$$\begin{array}{ll} \text{Atomic expressions} & a ::= c \mid x \mid \lambda x.k \mid (x, x) \mid xx \mid \pi_i x \mid (x \in \tau) ? x : x \\ \text{Canonical Forms} & \kappa ::= x \mid \text{bind } x = a \text{ in } \kappa \end{array}$$

Canonical forms, ranged over by κ , are *binding variables* (noted x , y , or z) possibly preceded by a list of definitions (from binding variables to atoms). Atoms are either a variable from a λ -abstraction (noted x , y , or z), or a constant, or a λ -abstraction whose body is a canonical form, or any other expression in which *all* proper sub-expressions are binding variables. An expression in canonical form without any free binding variable can be transformed into an expression of the source language using the unwinding operator $[\cdot]$ that basically inlines bindings: $[\text{bind } x = a \text{ in } \kappa] = [\kappa] \{ [a] / x \}$ (see Appendix E.1 for the full definition). The inverse direction, that is, producing from a source language expression a canonical form that unwinds to it, is straightforward (see Appendix E.2). However for each expression of the source language there are several canonical forms that unwind to it. For our algorithmic type system we need to associate each source language expression to a unique canonical form, as we define next.

3.1.2 Maximal Sharing Canonical Forms. We define a congruence on canonical forms and atoms:

DEFINITION 3.1 (CANONICAL EQUIVALENCE). We denote by \equiv_κ the smallest congruence on canonical forms and atoms that is closed by α -conversion and such that

$$\text{bind } x_1 = a_1 \text{ in bind } x_2 = a_2 \text{ in } \kappa \equiv_\kappa \text{bind } x_2 = a_2 \text{ in bind } x_1 = a_1 \text{ in } \kappa \quad x_1 \notin \text{fv}(a_2), x_2 \notin \text{fv}(a_1)$$

To infer types for the source language, we single out canonical forms satisfying three properties:

DEFINITION 3.2 (MSC FORMS). A maximal sharing canonical form (*abbreviated as MSC-form*) is (any canonical form α -equivalent to) a canonical form κ such that:

- (1) if $\text{bind } x_1 = a_1 \text{ in } \kappa_1$ and $\text{bind } x_2 = a_2 \text{ in } \kappa_2$ are distinct sub-expressions of κ , then $a_1 \not\equiv_\kappa a_2$
- (2) if $\lambda x. \kappa_1$ is a sub-expression of κ and $\text{bind } y = a \text{ in } \kappa_2$ a sub-expression of κ_1 , then $\text{fv}(a) \not\subseteq \text{fv}(\lambda x. \kappa_1)$
- (3) if $\text{bind } x = a \text{ in } \kappa'$ is a sub-expression of κ , then $x \in \text{fv}(\kappa')$.

MSC-forms are defined modulo α -conversion.⁷ The first condition states that distinct variables denote different definitions, that is, it enforces the *maximal sharing* of common sub-expressions. The second condition requires bindings to extrude λ -abstractions whenever possible. The third condition states that there is no useless bind (bound variables must occur in the body of the binds).

The key property of MSC-forms is that given an expression e of the source language, all its MSC-forms (i.e., all MSC-form whose unwinding is e) are equivalent:

PROPOSITION 3.3. If κ_1 and κ_2 are two MSC-forms and $[\kappa_1] \equiv_\alpha [\kappa_2]$, then $\kappa_1 \equiv_\kappa \kappa_2$.

We denote the unique MSC-form whose unwinding is e by $\text{MSC}(e)$. It is easy to transform a canonical form into a MSC-form that has the same unwinding. The reader can refer to Appendix E for a set of rewriting rules implementing this operation.

3.2 Algorithmic Typing Rules

MSC-forms tell us when to apply the $[\vee]$ rule: a term $\text{bind } x = a \text{ in } \kappa$ means (roughly) that it must be typed by applying the union rule to the expression $\kappa \{ a / x \}$. Putting an expression into its MSC-form to type it, thus corresponds to applying the $[\vee]$ rule on every occurrence of every subexpression of the original expression. This is a step toward a syntax-directed type system. However, there are still two issues to solve before obtaining an algorithmic type system: (i) rules $[\wedge]$, $[\text{INST}]$, and $[\leq]$ are still not syntax-directed, and (ii) rules $[\vee]$, $[\text{INST}]$, $[\rightarrow\text{I}]$, and $[\leq]$ are not analytic, meaning that some of their premises cannot be deduced just by looking at the conclusion: the $[\vee]$ rule requires guessing a type decomposition (i.e., the monomorphic type \mathbf{u} in the premises), the $[\text{INST}]$ rule

⁷For instance, both $\lambda x. \text{bind } x = x \text{ in bind } z = xy \text{ in bind } z' = zy \text{ in } z'$ and $\lambda x. \text{bind } x = x \text{ in bind } z = xy \text{ in } z$ are two distinct atoms that can occur in the same MSC-form, even though the atom xy appears in both: an α -renaming of x makes the first MSC-property hold.

requires guessing a substitution, the $[\rightarrow I]$ rule requires guessing the domain \mathbf{u} of the function, and the $[\leq]$ rule requires guessing the type t' to subsume to.

The issue of $[\text{INST}]$ and $[\leq]$ not being syntax directed can be solved by embedding them in some structural rules (in particular, in the rules for destructors). Moreover, as we will see later, the rules in which we embed $[\leq]$ can be made analytic by using some type operators. As for rule $[\wedge]$, making it syntax-directed, is trickier. Indeed, the usual approach of merging rules $[\rightarrow I]$ and $[\wedge]$ does not work here, since terms in MSC-forms may hoist a bind definition outside the function where they are used, causing rule $[\wedge]$ to be needed on a term that is not, syntactically, a λ -abstraction. Lastly, there is no easy way to guess the substitutions used by $[\text{INST}]$ rules, or the domain used in $[\rightarrow I]$ rules, or the decompositions performed by $[\vee]$ rules. To tackle these issues, our algorithmic type system will not only take a canonical form as input, but also an annotation that will (i) indicate when to apply an intersection, and (ii) indicate which type decomposition (for $[\vee]$ rules) and which type substitutions (for $[\text{INST}]$ rules) to use. Formally, our algorithmic system uses judgements of the form $\Gamma \vdash_{\mathcal{A}} [\kappa \mid \mathbb{k}] : t$ for a canonical form κ , and $\Gamma \vdash_{\mathcal{A}} [a \mid \mathbb{a}] : t$ for an atom a where \mathbb{k} and \mathbb{a} are respectively *form annotations* and *atom annotations* defined as follows:

Atom annots $\mathbb{a} ::= \emptyset \mid \lambda(\mathbf{u}, \mathbb{k}) \mid (\rho, \rho) \mid @(\Sigma, \Sigma) \mid \pi(\Sigma) \mid \mathbb{O}(\Sigma) \mid \in_1(\Sigma) \mid \in_2(\Sigma) \mid \wedge(\{\mathbb{a}, \dots, \mathbb{a}\})$
Form annots $\mathbb{k} ::= \rho \mid \text{keep } (\mathbb{a}, \{\mathbf{u}, \mathbb{k}\}, \dots, (\mathbf{u}, \mathbb{k}\}) \mid \text{skip } \mathbb{k} \mid \wedge(\{\mathbb{k}, \dots, \mathbb{k}\})$

where ρ ranges over *renamings* of polymorphic variables, that is, injective substitutions from \mathcal{V}_P to \mathcal{V}_P , and Σ ranges over instantiations, that is, sets of substitutions from \mathcal{V}_P to **Types**. We chose to keep annotations separate from the terms, instead of embedding them in the canonical forms, since in the latter case it would be more complicated to capture the tree structure of the derivations.

The algorithmic system is defined by the rules given in Appendix G. Below we comment the most interesting rules (we just omit the rules for constants, variables and two rules for type-cases). Essentially, there is one typing rule for each annotation, the only exception being the \emptyset annotation that is used both in the rule to type constants and in the two rules for variables.

$$[\rightarrow I\text{-ALG}] \frac{\Gamma, x : \mathbf{u} \vdash_{\mathcal{A}} [\kappa \mid \mathbb{k}] : t}{\Gamma \vdash_{\mathcal{A}} [\lambda x. \kappa \mid \lambda(\mathbf{u}, \mathbb{k})] : \mathbf{u} \rightarrow t}$$

To type the atom $\lambda x. \kappa$, the annotation $\lambda(\mathbf{u}, \mathbb{k})$ provides the domain \mathbf{u} of the function, and the annotation \mathbb{k} for its body.

$$[\rightarrow E\text{-ALG}] \frac{t_1 = \Gamma(x_1)\Sigma_1, \quad t_2 = \Gamma(x_2)\Sigma_2}{\Gamma \vdash_{\mathcal{A}} [x_1 x_2 \mid @(\Sigma_1, \Sigma_2)] : t_1 \circ t_2 \quad t_1 \leq \mathbb{O} \rightarrow \mathbb{1}, \quad t_2 \leq \text{dom}(t_1)}$$

To type an application one must apply an instantiation and a subsumption to both the type of the function and the type of the argument. Instantiations (i.e., Σ_1 and Σ_2) are sets of type substitutions; their application to a type t is defined as $t\Sigma \stackrel{\text{def}}{=} \bigwedge_{\sigma \in \Sigma} t\sigma$. Since they cannot be directly guessed, they are given by the annotation. Subsumption instead is embedded in two type operators. A first operator, $\text{dom}()$, computes the domain of the arrow and is used to check that the application is well-typed. A second type operator, \circ , computes the type of the result of the application. These type operators are defined as follows: $\text{dom}(t) \stackrel{\text{def}}{=} \max\{u \mid t \leq u \rightarrow \mathbb{1}\}$ and $t \circ s \stackrel{\text{def}}{=} \min\{u \mid t \leq s \rightarrow u\}$.

$$[\times E_1\text{-ALG}] \frac{t = \Gamma(x)\Sigma}{\Gamma \vdash_{\mathcal{A}} [\pi_1 x \mid \pi(\Sigma)] : \pi_1(t) \quad t \leq (\mathbb{1} \times \mathbb{1})} \quad [\times E_2\text{-ALG}] \frac{t = \Gamma(x)\Sigma}{\Gamma \vdash_{\mathcal{A}} [\pi_2 x \mid \pi(\Sigma)] : \pi_2(t) \quad t \leq (\mathbb{1} \times \mathbb{1})}$$

The rules for projections $[\times E_1\text{-ALG}]$ and $[\times E_2\text{-ALG}]$ follow the same idea as the rule for application $[\rightarrow E\text{-ALG}]$, with the use of two type operators $\pi_1(t) \stackrel{\text{def}}{=} \min\{u \mid t \leq u \times \mathbb{1}\}$ and $\pi_2(t) \stackrel{\text{def}}{=} \min\{u \mid t \leq \mathbb{1} \times u\}$. All these type operators can be effectively computed (cf. Appendix F).

$$[\times\text{-ALG}] \frac{}{\Gamma \vdash_{\mathcal{A}} [(x_1, x_2) \mid (\rho_1, \rho_2)] : t_1 \times t_2} \quad t_1 = \Gamma(x_1)\rho_1, \quad t_2 = \Gamma(x_2)\rho_2$$

To type a pair (x_1, x_2) it is not necessary to instantiate $\Gamma(x_1)$ or $\Gamma(x_2)$. However, to avoid unwanted correlations, it is necessary to rename the polymorphic type variables of its components. For instance, when typing the pair (x, x) with $x : \alpha \rightarrow \alpha$, it is better to type it with $(\alpha \rightarrow \alpha, \beta \rightarrow \beta)$ rather than $(\alpha \rightarrow \alpha, \alpha \rightarrow \alpha)$, since the former type has strictly more instances than the latter.

$$[\in_1\text{-ALG}] \frac{}{\Gamma \vdash_{\mathcal{A}} [(x \in \tau) ? x_1 : x_2 \mid \in_1(\Sigma)] : \Gamma(x)\Sigma \leq \tau}$$

To type type-cases, the annotation indicates which of the three rules must be applied (here $[\in_1]$) and how to instantiate the polymorphic type variables occurring in the type of the tested expression, so that it satisfies the side condition of the applied rule (see also $[\in_2\text{-ALG}]$ and $[\emptyset\text{-ALG}]$ in Appendix G).

$$[\text{BIND}_1\text{-ALG}] \frac{\Gamma \vdash_{\mathcal{A}} [\kappa \mid \mathbb{k}] : t}{\Gamma \vdash_{\mathcal{A}} [\text{bind } x = a \text{ in } \kappa \mid \text{skip } \mathbb{k}] : t} \quad x \notin \text{dom}(\Gamma)$$

In rule $[\text{BIND}_1\text{-ALG}]$ the annotation indicates to skip the definition of the current binding. This rule is used when the binding variable is not required for typing the body κ under the current context Γ . For instance, this is the case when x only appears in a branch of a typecase that cannot be taken under Γ . The side condition $x \notin \Gamma$ prevents a potential unsound name conflict between binding variables: as occurrences of x in κ denote the x binding variable that is being skipped, having the type of a former binding variable x in our environment when typing κ would be unsound.

$$[\text{BIND}_2\text{-ALG}] \frac{\Gamma \vdash_{\mathcal{A}} [a \mid \emptyset] : s \quad (\forall i \in I) \quad \Gamma, x : s \wedge \mathbf{u}_i \vdash_{\mathcal{A}} [\kappa \mid \mathbb{k}_i] : t_i}{\Gamma \vdash_{\mathcal{A}} [\text{bind } x = a \text{ in } \kappa \mid \text{keep } (\emptyset, \{\mathbf{u}_i, \mathbb{k}_i\}_{i \in I})] : \bigvee_{i \in I} t_i} \quad \bigvee_{i \in I} \mathbf{u}_i \approx \mathbb{1}$$

This rule tries to type the bound atom and then decomposes its type according to the annotation. This decomposition corresponds to an application of the $[\vee]$ rule of the declarative type system with the only difference that the type s of the atom is split in several summands by intersecting it with the various \mathbf{u}_i (instead of just two summands as in the rule $[\vee]$) whose union covers $\mathbb{1}$.

Finally, two annotations indicate when and how to apply rule $[\wedge]$ to atoms and canonical forms:

$$[\wedge\text{-ALG}] \frac{(\forall i \in I) \quad \Gamma \vdash_{\mathcal{A}} [a \mid \emptyset_i] : t_i}{\Gamma \vdash_{\mathcal{A}} [a \mid \wedge(\{\emptyset_i\}_{i \in I})] : \bigwedge_{i \in I} t_i} \quad I \neq \emptyset \quad [\wedge\text{-ALG}] \frac{(\forall i \in I) \quad \Gamma \vdash_{\mathcal{A}} [\kappa \mid \mathbb{k}_i] : t_i}{\Gamma \vdash_{\mathcal{A}} [\kappa \mid \wedge(\{\mathbb{k}_i\}_{i \in I})] : \bigwedge_{i \in I} t_i} \quad I \neq \emptyset$$

An expression e is typable if and only if its unique (modulo \equiv_{κ}) MSC-form is typable, too:

THEOREM 3.4 (SOUNDNESS AND COMPLETENESS). *For every term e of the source language*

$$\begin{aligned} \vdash e : t &\Rightarrow \exists \mathbb{k} \vdash_{\mathcal{A}} [\text{MSC}(e) \mid \mathbb{k}] : t' \leq t && \text{(completeness)} \\ \vdash e : t &\Leftarrow \vdash_{\mathcal{A}} [\text{MSC}(e) \mid \mathbb{k}] : t && \text{(soundness)} \end{aligned}$$

It is easy to generate the unique MSC-form associated to a source language expression e (cf. Appendix E). Theorem 3.4 states that this MSC-form is typable if and only if e is: we reduced the problem of typing e to the one of finding an annotation that makes the unique MSC-form of e typeable with the algorithmic type system.⁸ Figure 3 gives an example of an MSC-form and two possible annotations for it. The term “ $\lambda x. (f x \in \text{Int}) ? g(f x) : x$ ” (where $f : \forall \alpha. \alpha \rightarrow \alpha$ and

⁸As stated by Theorem 3.4 the transformation of an expression into its MSC-form preserves typing. However, intuitively, it does not preserve the reduction semantics, since bindings regroup different occurrences of some sub-expression that in the original expression might be evaluated at different stages of the reduction or not evaluated at all. We said “intuitively” because no operational semantics is defined for canonical forms (this was already the case for [Castagna et al. 2022b]).

$\Gamma = \{f : \alpha \rightarrow \alpha, g : \text{Int} \rightarrow \text{Int}\}$		
<pre> 1 bind z = 2 3 $\lambda x.$ 4 bind x = x in 5 bind u = f x in 6 7 bind v = g u in 8 bind w = 9 (u ∈ Int) ? v : x 10 in w 11 12 13 in z </pre>	<pre> keep($\lambda(\beta,$ keep($\emptyset, \{\mathbb{1},$ keep($\@(\{\{\alpha \rightsquigarrow \beta\}, \{\emptyset\}),$ {Int, keep($\@(\{\emptyset, \{\emptyset\}, \{\mathbb{1},$ keep($\in_1(\{\emptyset\}, \{\{\mathbb{1}, \emptyset\}\})$), ($\neg$Int, skip(keep($\in_2(\{\emptyset\}, \{\{\mathbb{1}, \emptyset\}\})$))))))))))))))))))))))) </pre>	<pre> keep($\wedge(\{$ $\lambda(\beta \setminus \text{Int},$ keep($\emptyset, \{\mathbb{1},$ keep($\@(\{\{\alpha \rightsquigarrow \text{Int}\},$ keep($\@(\{\{\alpha \rightsquigarrow \beta \setminus \text{Int}\},$ {Int, keep($\@(\{\emptyset, \{\emptyset\}, \{\mathbb{1},$ skip(keep($\in_2(\{\emptyset\},$ $\{\{\mathbb{1}, \emptyset\}\})$))))))))))))))))))))))))))) </pre>
FINAL TYPE:	$\beta \rightarrow \beta \vee \text{Int}$	$(\text{Int} \rightarrow \text{Int}) \wedge ((\beta \setminus \text{Int}) \rightarrow (\beta \setminus \text{Int}))$

Fig. 3. MSC-form and two annotations for $\lambda x. (fx \in \text{Int}) ? g(fx) : x$

$g : \text{Int} \rightarrow \text{Int}$ is put in MSC-form (on the left). In the first annotation, the function is typed with a single λ annotation (line 3). The interesting part is the annotation of the binding for u (line 5): the corresponding keep annotation represents an application of the union elimination rule on the occurrences of the expression fx whose type β is split into $\beta \wedge \text{Int}$ (line 6) and $\beta \setminus \text{Int}$ (line 9). Each subcase is annotated accordingly. Notice in the second subcase that the annotation for v is skip (line 10) which indicates that this particular variable must not be used (as $g(fx)$ cannot be typed since in the “else” branch, fx has type $\neg \text{Int}$). A different annotation, yielding a better type, is the one on the right. This intersection annotation (line 2) separates the domain of the λ -abstraction into two cases, each typed independently, yielding for the whole function an intersection type.

The example in Figure 3 also shows why the condition of maximal sharing for our forms is necessary, not only for their uniqueness, but also for the completeness of the algorithmic system: if the two occurrences of fx in “ $\lambda x. (fx \in \text{Int}) ? g(fx) : x$ ” were not bound by the same variable (as in the leftmost column of line 5 in Figure 3), viz., if the sharing were not maximal, then it would not be possible to deduce that $g(fx)$ is well typed: g expects an integer, but without maximal sharing it is not possible to deduce that the occurrence of fx in the first branch is indeed of type Int .

The problem of inferring an annotation for an MSC-form as the above—in particular the rightmost (more precise) annotation in Figure 3—is tackled in the next section.

4 RECONSTRUCTION

This section describes an algorithm to find an annotation for an expression in MSC-form, such that the pair expression and annotation is typable in the algorithmic system. Though this algorithm is not complete, it is sound and terminating (see Section 4.4 for the formal statements and a discussion about incompleteness). Experimental results are presented in Section 5.

The annotation reconstruction algorithm is composed of two systems of deduction rules: the *main reconstruction algorithm* (Section 4.2) which produces intermediate annotations containing information about the domains of λ -abstractions and the type decompositions to use in bindings, and the *auxiliary reconstruction algorithm* (Section 4.3) which converts these intermediate annotations into annotations for the algorithmic type system, by computing instantiations Σ for destructors.

4.1 The Tallying Algorithm

One key ingredient used by the reconstruction algorithm is the *tallying* algorithm. Roughly, *tallying* is the equivalent of the *unification* used in algorithm \mathcal{W} [Damas and Milner 1982], but for a type system with subtyping. The tallying algorithm was introduced by Castagna et al. [2015] to solve the following problem: given a set of pairs $\{(t_i, t'_i)\}_{i \in I}$ and a set of type variables Δ representing the monomorphic type variables, find all substitutions σ whose domain is disjoint from Δ (noted $\sigma \# \Delta$) and that satisfy $\forall i \in I. t_i \sigma \leq t'_i \sigma$. Castagna et al. [2015] show that this problem is decidable and give an algorithm to characterize all solutions. As for unification, for each instance of the tallying problem there is either no solution or several substitutions each of which is a solution of the problem. The difference is that while with unification all solutions are characterized by a principal substitution, with tallying they are characterized by a principal finite set of substitutions.⁹ More precisely, all substitutions that are solutions to a tallying instance are characterized by a *principal* set Σ of substitutions, such that every $\sigma \in \Sigma$ is a solution, and for any solution σ , we have $\exists \sigma_1 \in \Sigma. \exists \sigma_2. \sigma_2 \# \Delta$ and $\sigma \approx \sigma_2 \circ \sigma_1$, where \circ denotes the composition of substitutions and \approx is pointwise type equivalence.

In this work, all tallying instances use a single constraint, and we will note $\text{tally}(\{t_1 \dot{\leq} t_2\})$ the set of substitutions Σ characterizing all the solutions of the tallying instance $\{(t_1, t_2)\}$, where $\Delta = \mathcal{V}_M$ and thus $\forall \sigma \in \Sigma. \text{dom}(\sigma) \subseteq \mathcal{V}_P$ (we use the symbol $\dot{\leq}$, rather than \leq to stress that it denotes a constraint to be solved, rather than a pair in the subtyping relation).

The tallying function $\text{tally}()$ finds substitutions for polymorphic type variables, but in order to infer the domain of λ -abstractions, we may need to find substitutions for monomorphic type variables. We thus introduce an additional tallying function, $\text{tally_infer}(\{t_1 \dot{\leq} t_2\})$:

DEFINITION 4.1. Let $\sigma|_X$ denote the restriction of the substitution σ to the domain X . We define

$$\text{tally_infer}(\{t_1 \dot{\leq} t_2\}) = \{(\sigma \circ \sigma' \circ \phi)|_{\mathcal{V}_M} \mid \sigma' \in \text{tally}(\{\text{fresh}(t_1)\phi \dot{\leq} \text{fresh}(t_2)\phi\})\}$$

where $\text{fresh}(t)$ denotes the type t where polymorphic type variables have been substituted by fresh ones; ϕ is a renaming from $(\text{vars}(t_1) \cup \text{vars}(t_2)) \cap \mathcal{V}_M$ to fresh polymorphic variables; and σ is a substitution mapping polymorphic variables appearing in the image of $\sigma' \circ \phi$ to fresh monomorphic variables.

In a nutshell, polymorphic type variables in t_1 and t_2 are refreshed in order to decorrelate them, and monomorphic type variables are generalized using ϕ so that $\text{tally}()$ is allowed to find solutions for them. Each solution σ' is composed with ϕ in order to restore the connection with the initial monomorphic type variables, and the polymorphic type variables in the image of the resulting substitution are transformed into monomorphic ones by composing σ with it. Finally, the substitution is restricted to \mathcal{V}_M (i.e., to the domain of ϕ).

For example, an instance such as $\text{tally_infer}(\{\text{Int} \wedge \alpha \rightarrow \text{Int} \wedge \alpha \dot{\leq} \beta \rightarrow \alpha\})$ can be generated during reconstruction, when a function of type $\text{Int} \wedge \alpha \rightarrow \text{Int} \wedge \alpha$ is applied to an argument of type β , but the α on the right-hand of $\dot{\leq}$ is unrelated to the one on the left-hand side. Decorrelating them yields a unique solution $\{\beta \rightsquigarrow \beta' \wedge \text{Int}\}$, that is, β must be substituted by $\beta' \wedge \text{Int}$ in our context for the application to be typeable.

4.2 Main Reconstruction Algorithm

The *main reconstruction algorithm*, defined in this section, infers the domains of λ -abstractions and the decompositions of types into disjoint unions to use for bindings. It works by successively

⁹This is due to the presence of the empty type. For instance, the principal solution of unifying $\alpha \times \beta$ with $s \times t$ is the substitution $\{\alpha \rightsquigarrow s, \beta \rightsquigarrow t\}$, while all substitutions that make the former type become a subtype of the latter are characterized by a set containing three distinct substitutions: $\{\alpha \rightsquigarrow \emptyset\}$, $\{\beta \rightsquigarrow \emptyset\}$, and $\{\alpha \rightsquigarrow s, \beta \rightsquigarrow t\}$.

refining *intermediate annotations* defined below. These intermediate annotations store information about the domains of λ -abstractions and the decompositions of bindings. However, the instantiations Σ used to type destructors (i.e., applications, projections, and typecases) in the algorithmic type system are not stored in intermediate annotations, because they might get invalidated as the reconstruction progresses: when new information is found about the domain of a lambda or the decomposition of a binding, the algorithm will retype some intermediate definitions of the MSC-form, thus invalidating the instantiations Σ of later definitions. Thus, these instantiations Σ will be recomputed whenever needed, using the auxiliary system (Section 4.3) that converts intermediate annotations into annotations for the algorithmic type system.

Atom and form *intermediate annotations* are defined by the grammar below:

Split annotations	$S ::= \{(\mathbf{u}, \mathcal{K}), \dots, (\mathbf{u}, \mathcal{K})\}$
Atom intermediate annot.	$\mathcal{A} ::= \text{infer} \mid \text{untyp} \mid \text{typ} \mid \wedge(\{\mathcal{A}, \dots, \mathcal{A}\}, \{\mathcal{A}, \dots, \mathcal{A}\})$ $\mid \in_1 \mid \in_2 \mid \lambda(\mathbf{u}, \mathcal{K})$
Form intermediate annot.	$\mathcal{K} ::= \text{infer} \mid \text{untyp} \mid \text{typ} \mid \wedge(\{\mathcal{K}, \dots, \mathcal{K}\}, \{\mathcal{K}, \dots, \mathcal{K}\})$ $\mid \text{try-skip}(\mathcal{K}) \mid \text{try-keep}(\mathcal{A}, \mathcal{K}, \mathcal{K})$ $\mid \text{propagate}(\mathcal{A}, \mathbb{F}, \mathcal{S}, \mathcal{S}) \mid \text{skip}(\mathcal{K}) \mid \text{keep}(\mathcal{A}, \mathcal{S}, \mathcal{S})$

In the following, we use the metavariable η to range over both atoms and expressions (i.e., $\eta ::= a \mid \kappa$). Similarly, the metavariable \mathfrak{h} ranges over atom annotations \mathfrak{a} and form annotations \mathfrak{k} (i.e., $\mathfrak{h} ::= \mathfrak{a} \mid \mathfrak{k}$); while the metavariable \mathcal{H} ranges over atom intermediate annotations \mathcal{A} and form intermediate annotations \mathcal{K} (i.e. $\mathcal{H} ::= \mathcal{A} \mid \mathcal{K}$).

Let ψ range over *monomorphic substitution*, that is, substitutions from \mathcal{V}_M to monomorphic types, and Ψ range over finite sets of monomorphic substitutions ($\Psi ::= \{\psi, \dots, \psi\}$). The main reconstruction algorithm is presented as a deduction rule system, for judgments of the form $\Gamma \vdash_{\mathcal{R}}^* \langle \eta \mid \mathcal{H} \rangle \Rightarrow \mathbb{R}$, where \mathbb{R} is a result defined as follows:

Result $\mathbb{R} ::= \text{Ok}(\mathcal{H}) \mid \text{Fail} \mid \text{Split}(\Gamma, \mathcal{H}, \mathcal{H}) \mid \text{Subst}(\Psi, \mathcal{H}, \mathcal{H}) \mid \text{Var}(x, \mathcal{H}, \mathcal{H})$

Let us see what each result for $\Gamma \vdash_{\mathcal{R}}^* \langle \eta \mid \mathcal{H} \rangle$ means:

Ok(\mathcal{H}'): the reconstruction was successful and η can be typed by the algorithmic type system using the annotation \mathcal{H}' (after converting it into an annotation \mathfrak{h} using the auxiliary reconstruction system). This result is terminal (i.e., it is a definitive answer that cannot be further refined).

Fail: the reconstruction has failed. The algorithm was not able to find an annotation that makes η typable with the algorithmic system. This result is terminal.

Subst($\Psi, \mathcal{H}_1, \mathcal{H}_2$): the reconstruction found a set of substitutions Ψ that if applied to Γ may make η typable. In practice, for each substitution $\psi \in \Psi$, the reconstruction will be called again on the environment $\Gamma\psi$ and annotation $\mathcal{H}_1\psi$. However, this does not necessarily mean that the reconstruction will fail on the current environment Γ : η might still be typeable but with a less precise type (e.g., it could yield an arrow type with a smaller domain). Thus, this *default* case which does not instantiate Γ is also explored, using the annotation \mathcal{H}_2 instead of \mathcal{H}_1 .

Split($\Gamma', \mathcal{H}_1, \mathcal{H}_2$): the reconstruction found some splits for the variables in $\text{dom}(\Gamma')$ that if applied to Γ may make η typable. In practice, the system generates several new environments: one is obtained by (pointwise) intersecting Γ with Γ' and then it is used to retype η with the annotation \mathcal{H}_1 ; the others are obtained by intersecting Γ with all the possible pointwise negations of Γ' and then they are used to retype η with the annotation \mathcal{H}_2 .

Var($x, \mathcal{H}_1, \mathcal{H}_2$): the reconstruction found that in order to type η , the definition of the bind-abstracted variable x should be typed. Any branch that successfully types it continues with the annotation \mathcal{H}_1 , otherwise it continues with the annotation \mathcal{H}_2 .

Initially, any form or atom η is annotated with `infer`, and this annotation is then refined until it yields a terminal result (i.e., either `Ok()` or `Fail`). The rules below are presented by decreasing priority (i.e., the first rule that applies is used). Some rules have been omitted for concision, but the reader can find the full reconstruction system in Appendix H.1.

There are two different forms of judgments: $\Gamma \vdash_{\mathcal{R}}^* \langle \eta \mid \mathcal{H} \rangle \Rightarrow \mathbb{R}$ and $\Gamma \vdash_{\mathcal{R}} \langle \eta \mid \mathcal{H} \rangle \Rightarrow \mathbb{R}$. We first define rules for the judgment $\vdash_{\mathcal{R}}$ for every canonical form and atom. The results of these judgments are not necessarily terminal and, therefore, it may be necessary to call the reconstruction again in order to refine them. This is the purpose of $\vdash_{\mathcal{R}}^*$ judgments which call repetitively $\vdash_{\mathcal{R}}$ judgments when relevant, so that in the end we get a terminal result. Let us first focus on $\vdash_{\mathcal{R}}$ judgments.

$$[\text{OK}] \frac{}{\Gamma \vdash_{\mathcal{R}} \langle \eta \mid \text{typ} \rangle \Rightarrow \text{Ok}(\text{typ})} \quad [\text{FAIL}] \frac{}{\Gamma \vdash_{\mathcal{R}} \langle \eta \mid \text{untyp} \rangle \Rightarrow \text{Fail}}$$

If a canonical form or atom η is annotated with `typ`, then reconstruction is finished for η , and it is typeable in the current context Γ . The annotation `typ` is never used on λ -abstractions and bindings because the system needs to store more information for them. Likewise, if a form or atom η is annotated with `untyp`, then reconstruction is finished for η by failing in the current context.

$$[\text{AxOK}] \frac{x \in \text{dom}(\Gamma)}{\Gamma \vdash_{\mathcal{R}} \langle x \mid \text{infer} \rangle \Rightarrow \text{Ok}(\text{typ})} \quad [\text{AxFAIL}] \frac{}{\Gamma \vdash_{\mathcal{R}} \langle x \mid \text{infer} \rangle \Rightarrow \text{Fail}}$$

If a λ -abstracted variable x is in the environment, then it is typeable and thus the algorithm returns `Ok(typ)`. Otherwise, x is undefined and `Fail` is returned.

$$[\text{APPVAR}_i] \frac{x_i \notin \text{dom}(\Gamma)}{\Gamma \vdash_{\mathcal{R}} \langle x_1 x_2 \mid \text{infer} \rangle \Rightarrow \text{Var}(x_i, \text{infer}, \text{untyp})}$$

To type the application $x_1 x_2$, we must first ensure that $\{x_1, x_2\} \subseteq \text{dom}(\Gamma)$. If it is not the case, then the two rules $[\text{APPVAR}_i]$ (for $i = 1, 2$) try to remedy it by returning `Var(x_i , infer, untyp)`, which is the result that asks the system to try to type the atom bound to x_i for $x_i \notin \text{dom}(\Gamma)$. If the attempt is successful, then the algorithm will continue the reconstruction for the application with the annotation `infer` and $x_i \in \text{dom}(\Gamma)$, otherwise it will continue with the annotation `untyp` making the reconstruction fail on this application.

$$[\text{APPINFER}] \frac{\Psi = \text{tally_infer}(\{\Gamma(x_1) \dot{\leftarrow} \Gamma(x_2) \rightarrow \alpha\})}{\Gamma \vdash_{\mathcal{R}} \langle x_1 x_2 \mid \text{infer} \rangle \Rightarrow \text{Subst}(\Psi, \text{typ}, \text{untyp})} \quad \alpha \in \mathcal{V}_p \text{ fresh}$$

If $\{x_1, x_2\} \subseteq \text{dom}(\Gamma)$, then the rule $[\text{APPINFER}]$ tries to find all instances of the current context in which the application $x_1 x_2$ is typeable, by subsuming $\Gamma(x_1)$ (the type of the function) to $\Gamma(x_2) \rightarrow \alpha$ (a function type whose domain is the type of the argument). For that, it calls the tallying algorithm which returns a set of substitutions Ψ . Then, `Subst(Ψ , typ, untyp)` is returned, meaning that this application should be typeable under every instance $\Gamma\psi$ of the current context Γ (with $\psi \in \Psi$). The default case (i.e., when the current context is unchanged, for example, when $\Psi = \emptyset$) cannot be typed, so it is annotated with `untyp` (see rule $[\text{ITERATE}_2]$ later on). The rules for pairs are similar and have been omitted.

$$[\text{CASESPLIT}] \frac{\Gamma(x) \not\leq \tau \quad \Gamma(x) \not\leq \neg\tau}{\Gamma \vdash_{\mathcal{R}} \langle (x \in \tau) ? x_1 : x_2 \mid \text{infer} \rangle \Rightarrow \text{Split}(\{(x : \tau)\}, \text{infer}, \text{infer})}$$

The key rule for type-cases is $[\text{CASESPLIT}]$, corresponding to the case where x is in Γ , but with a type that does not allow the selection of a specific branch. Thus, we need to partition the type of x in two, one part being a subtype of τ and the other a subtype of $\neg\tau$. This is achieved by returning

$\text{Split}(\{(x : \tau)\}, \text{infer}, \text{infer})$: this result is backtracked up to the binding of x , where it will split the associated type, accordingly.

$$\begin{aligned} [\text{CASETHEN}] & \frac{\Gamma(x) \leq \tau \quad \Psi = \text{tally_infer}(\{\Gamma(x) \dot{\leq} 0\})}{\Gamma \vdash_{\mathcal{R}} \langle (x \in \tau) ? x_1 : x_2 \mid \text{infer} \rangle \Rightarrow \text{Subst}(\Psi, \text{typ}, \epsilon_1)} \\ [\text{CASEELSE}] & \frac{\Gamma(x) \leq \neg\tau \quad \Psi = \text{tally_infer}(\{\Gamma(x) \dot{\leq} 0\})}{\Gamma \vdash_{\mathcal{R}} \langle (x \in \tau) ? x_1 : x_2 \mid \text{infer} \rangle \Rightarrow \text{Subst}(\Psi, \text{typ}, \epsilon_2)} \\ [\text{CASEVAR}_i] & \frac{x_i \notin \text{dom}(\Gamma)}{\Gamma \vdash_{\mathcal{R}} \langle (x \in \tau) ? x_1 : x_2 \mid \epsilon_i \rangle \Rightarrow \text{Var}(x_i, \text{typ}, \text{untyp})} \end{aligned}$$

When the type of x allows the selection of a branch, then either the rule $[\text{CASETHEN}]$ or the rule $[\text{CASEELSE}]$ applies. If we are in the case of $[\text{CASETHEN}]$, that is $\Gamma(x) \leq \tau$, then we have to determine whether we will apply the algorithmic rule $[\text{0-ALG}]$ or the algorithmic rule $[\epsilon_1\text{-ALG}]$. To determine it, the $[\text{CASETHEN}]$ rule calls $\text{tally_infer}(\{\Gamma(x) \dot{\leq} 0\})$ which returns the set of contexts $\Gamma\psi$ (for $\psi \in \Psi$) under which the algorithmic rule $[\text{0-ALG}]$ is to be applied, that is, the contexts under which the tested expression x has an empty type. The default case, corresponding to the case in which the type of $\Gamma(x)$ is not guaranteed to be empty and, thus, in which the algorithmic rule $[\epsilon_1\text{-ALG}]$ must be applied, is annotated with ϵ_1 . This annotation is handled by the rule $[\text{CASEVAR}_1]$ which forces the system to type x_1 , the binding variable associated to the first branch. The case for $[\text{CASEELSE}]$ and $[\text{CASEVAR}_2]$ is analogous.

We omitted the remaining rules for type-cases since they are straightforward: the rule for $x \notin \text{dom}(\Gamma)$, which triggers a $\text{Var}(x, \text{infer}, \text{untyp})$ result; two rules similar to $[\text{CASEVAR}_i]$, but where $x_i \in \text{dom}(\Gamma)$, which simply return $\text{Ok}(\text{typ})$.

$$\begin{aligned} [\text{LAMBDAINFERR}] & \frac{\Gamma \vdash_{\mathcal{R}} \langle \lambda x. \kappa \mid \lambda(\alpha, \text{infer}) \rangle \Rightarrow \mathbb{R}}{\Gamma \vdash_{\mathcal{R}} \langle \lambda x. \kappa \mid \text{infer} \rangle \Rightarrow \mathbb{R}} \quad \alpha \in \mathcal{V}_M \text{ fresh} \\ [\text{LAMBDA}] & \frac{\Gamma, x : \mathbf{u} \vdash_{\mathcal{R}}^* \langle \kappa \mid \mathcal{K} \rangle \Rightarrow \mathbb{R}}{\Gamma \vdash_{\mathcal{R}} \langle \lambda x. \kappa \mid \lambda(\mathbf{u}, \mathcal{K}) \rangle \Rightarrow \text{map}(X \mapsto \lambda(\mathbf{u}, X), \mathbb{R})} \end{aligned}$$

The rules for λ -abstractions mimic algorithm \mathcal{W} . Rule $[\text{LAMBDAINFERR}]$ transforms the initial infer annotation into a $\lambda(\alpha, \text{infer})$ annotation. As in \mathcal{W} , λ -abstracted variables are initially given a fresh type variable, which will then be substituted as needed while reconstructing the type of the body; here we use a fresh monomorphic variable, but $\text{tally_infer}()$ will transform it into a polymorphic—thus, instantiable—one, just for the reconstruction in the body. Rule $[\text{LAMBDA}]$ adds the λ -abstracted variable to the environment with the type specified in the annotation, recursively calls reconstruction on the body, and reestablishes the variable type annotation on the result. The notation $\text{map}(X \mapsto f(X), \mathbb{R})$ denotes the result \mathbb{R} where f has been applied to every annotation X .

$$[\text{BINDINFERR}] \frac{\Gamma \vdash_{\mathcal{R}} \langle \text{bind } x = a \text{ in } \kappa \mid \text{try-skip}(\text{infer}) \rangle \Rightarrow \mathbb{R}}{\Gamma \vdash_{\mathcal{R}} \langle \text{bind } x = a \text{ in } \kappa \mid \text{infer} \rangle \Rightarrow \mathbb{R}}$$

The $[\text{BINDINFERR}]$ rule transforms an initial infer annotation into a $\text{try-skip}(\text{infer})$ annotation which skips the binding and annotates the body κ with infer . We do not try to type the definition of a binding until it is actually used, because its variable might appear only in unreachable positions (e.g., in an unreachable branch of a type-case). In other words, we implement a lazy typing discipline for bind -abstracted variables. If the variable is used at some point, then an attempt to type it will

be initiated by the [BINDTRYSKIP₁] rule below:

$$[\text{BINDTRYSKIP}_1] \frac{\Gamma \vdash_{\mathcal{R}}^* \langle \kappa \mid \mathcal{K} \rangle \Rightarrow \text{Var}(x, \mathcal{K}_1, \mathcal{K}_2) \quad \Gamma \vdash_{\mathcal{R}} \langle \text{bind } x = a \text{ in } \kappa \mid \text{try-keep}(\text{infer}, \mathcal{K}_1, \mathcal{K}_2) \rangle \Rightarrow \mathbb{R}}{\Gamma \vdash_{\mathcal{R}} \langle \text{bind } x = a \text{ in } \kappa \mid \text{try-skip}(\mathcal{K}) \rangle \Rightarrow \mathbb{R}}$$

This rule tries to type the body of the binding, starting with the annotation \mathcal{K} (initially, infer). If the result is $\text{Var}(x, \mathcal{K}_1, \mathcal{K}_2)$, then it means that the current binding is used in the body κ and, thus, the system should try to type it. Consequently, the annotation for the current binding is changed into a $\text{try-keep}(\text{infer}, \mathcal{K}_1, \mathcal{K}_2)$ so that, at the next iteration, its definition will be reconstructed.

If typing the body of the binding yields a result different from $\text{Var}(x, \mathcal{K}_1, \mathcal{K}_2)$, then this result is just propagated as in [LAMBDA] (the corresponding rules have been omitted).

$$[\text{BINDTRYKEEP}_1] \frac{\Gamma \vdash_{\mathcal{R}}^* \langle a \mid \mathcal{A} \rangle \Rightarrow \text{Ok}(\mathcal{A}') \quad \Gamma \vdash_{\mathcal{R}} \langle \text{bind } x = a \text{ in } \kappa \mid \text{keep}(\mathcal{A}', \{(\perp, \mathcal{K}_1)\}, \emptyset) \rangle \Rightarrow \mathbb{R}}{\Gamma \vdash_{\mathcal{R}} \langle \text{bind } x = a \text{ in } \kappa \mid \text{try-keep}(\mathcal{A}, \mathcal{K}_1, \mathcal{K}_2) \rangle \Rightarrow \mathbb{R}}$$

$$[\text{BINDTRYKEEP}_2] \frac{\Gamma \vdash_{\mathcal{R}}^* \langle a \mid \mathcal{A} \rangle \Rightarrow \text{Fail} \quad \Gamma \vdash_{\mathcal{R}} \langle \text{bind } x = a \text{ in } \kappa \mid \text{skip}(\mathcal{K}_2) \rangle \Rightarrow \mathbb{R}}{\Gamma \vdash_{\mathcal{R}} \langle \text{bind } x = a \text{ in } \kappa \mid \text{try-keep}(\mathcal{A}, \mathcal{K}_1, \mathcal{K}_2) \rangle \Rightarrow \mathbb{R}}$$

As expected, if the current annotation for the binding is a $\text{try-keep}(\mathcal{A}, \mathcal{K}_1, \mathcal{K}_2)$, then the system tries to reconstruct the annotation for the definition. If it succeeds, then it becomes possible to type the definition and to continue the reconstruction of the body using \mathcal{K}_1 . This is what [BINDTRYKEEP₁] does by changing the current annotation to $\text{keep}(\mathcal{A}', \{(\perp, \mathcal{K}_1)\}, \emptyset)$ (more details below). If the reconstruction of the definition fails (rule [BINDTRYKEEP₂]), then we have no choice but to skip this definition and use the default annotation \mathcal{K}_2 to type the body.

In an annotation $\text{keep}(\mathcal{A}, \mathcal{S}, \mathcal{S}')$ for the binding of a variable x , \mathcal{A} is the annotation for typing the definition of x , while the two other arguments describe the type decomposition to use for x and, for each part of the decomposition, the annotation to use for the body. More precisely, \mathcal{S} contains the parts of the type decomposition that have yet to be explored, and \mathcal{S}' contains the parts that have already been fully explored. In particular, the annotation $\text{keep}(\mathcal{A}', \{(\perp, \mathcal{K}_1)\}, \emptyset)$ used in rule [BINDTRYKEEP₁] means that the type of the definition does not need to be partitioned: there is only one part, covering \perp , associated with an annotation \mathcal{K}_1 for typing the body.

$$[\text{BINDOK}] \frac{}{\Gamma \vdash_{\mathcal{R}} \langle \text{bind } x = a \text{ in } \kappa \mid \text{keep}(\mathcal{A}, \emptyset, \mathcal{S}) \rangle \Rightarrow \text{Ok}(\text{keep}(\mathcal{A}, \emptyset, \mathcal{S}))}$$

If all the parts of the type decomposition have already been explored (i.e., \emptyset in the annotation in the rule above), then the reconstruction is successful. Otherwise, the following rules are applied:

$$[\text{BINDKEEP}_1] \frac{\Gamma \vdash_{\mathcal{P}} \langle a \mid \mathcal{A} \rangle \Rightarrow \emptyset \quad \Gamma \vdash_{\mathcal{A}} [a \mid \emptyset] : s \quad \Gamma, x : s \wedge \mathbf{u} \vdash_{\mathcal{R}}^* \langle \kappa \mid \mathcal{K} \rangle \Rightarrow \text{Ok}(\mathcal{K}')}{\Gamma \vdash_{\mathcal{R}} \langle \text{bind } x = a \text{ in } \kappa \mid \text{keep}(\mathcal{A}, \mathcal{S}, \{(\mathbf{u}, \mathcal{K}')\} \cup \mathcal{S}') \rangle \Rightarrow \mathbb{R}}$$

$$\Gamma \vdash_{\mathcal{R}} \langle \text{bind } x = a \text{ in } \kappa \mid \text{keep}(\mathcal{A}, \{(\mathbf{u}, \mathcal{K})\} \cup \mathcal{S}, \mathcal{S}') \rangle \Rightarrow \mathbb{R}$$

$$[\text{BINDKEEP}_2] \frac{\Gamma \vdash_{\mathcal{P}} \langle a \mid \mathcal{A} \rangle \Rightarrow \emptyset \quad \Gamma \vdash_{\mathcal{A}} [a \mid \emptyset] : s \quad \Gamma, x : s \wedge \mathbf{u} \vdash_{\mathcal{R}}^* \langle \kappa \mid \mathcal{K} \rangle \Rightarrow \text{Split}(\Gamma', \mathcal{K}_1, \mathcal{K}_2)}{x \in \text{dom}(\Gamma') \quad \Gamma \vdash_{\mathcal{E}} (a : \neg(\mathbf{u} \wedge \Gamma'(x))) \Rightarrow \mathbb{F}_1 \quad \Gamma \vdash_{\mathcal{E}} (a : \neg(\mathbf{u} \searrow \Gamma'(x))) \Rightarrow \mathbb{F}_2}$$

$$\Gamma \vdash_{\mathcal{R}} \langle \text{bind } x = a \text{ in } \kappa \mid \text{keep}(\mathcal{A}, \{(\mathbf{u}, \mathcal{K})\} \cup \mathcal{S}, \mathcal{S}') \rangle \Rightarrow \text{Split}(\Gamma' \setminus x, \mathcal{K}'_1, \mathcal{K}'_2)$$

with, in the last rule, $\mathcal{K}'_1 = \text{propagate}(\mathcal{A}, \mathbb{F}_1 \cup \mathbb{F}_2, \{(\mathbf{u} \wedge \Gamma'(x), \mathcal{K}_1), (\mathbf{u} \searrow \Gamma'(x), \mathcal{K}_2)\} \cup \mathcal{S}, \mathcal{S}')$ and $\mathcal{K}'_2 = \text{keep}(\mathcal{A}, \{(\mathbf{u}, \mathcal{K}_2)\} \cup \mathcal{S}, \mathcal{S}')$

In both rules, the definition of the binding is typed using the annotation \mathcal{A} . For that, it is first converted into an annotation \mathfrak{a} of the algorithmic type system, using the deduction rules for the judgment $\Gamma \vdash_{\mathcal{P}} \langle a \mid \mathcal{A} \rangle \Rightarrow \mathfrak{a}$, defined in Section 4.3. Then, the type s obtained for the definition is intersected with one of the parts of the type decomposition, according to the second argument of the `keep()` annotation (i.e., $\{\mathbf{u}, \mathcal{K}\} \cup \mathcal{S}$ in both rules), and the corresponding annotation for the body is reconstructed recursively. Note that, since split annotations are sets, then the order in which the parts are explored is arbitrary.

The rule $[\text{BINDKEEP}_1]$ for an annotation `keep` $(\mathcal{A}, \mathcal{S}, \mathcal{S}')$ is responsible for moving a branch from \mathcal{S} to \mathcal{S}' when the result for the branch is `Ok()`. If instead the reconstruction of the body requires to further split the type of x , then the rule $[\text{BINDKEEP}_2]$ splits the current branch into two branches. However, before exploring these two branches, some information about the split needs to be propagated, to ensure that when a split is explored, it is under a context as precise as possible.

Let us explain this by an example. Assume we have a polymorphic primitive function `id` of type $\alpha \rightarrow \alpha$ and an initial environment $\Gamma = \{x : \text{Bool}\}$. We want to type the following canonical form, and deduce for it the type `True` (since x and y are always bound to the same value):

$$\text{bind } x = x \text{ in bind } y = \text{id } x \text{ in bind } z = (y \in \text{True}) ? x : \text{true in } z$$

At some point, the partition associated to y will change from $\{\mathbb{1}\}$ to $\{\text{True}, \neg\text{True}\}$ because of the type-case (rule $[\text{CASESPLIT}]$). However, if the case corresponding to $(y : \text{True})$ is immediately explored, it will yield for the body the type `Bool`, because x still has the type `Bool` in the environment. In order to obtain the more precise type `True`, we must deduce, before exploring the case $(y : \text{True})$, that when `id x` (the definition of y) has type `True`, then x also has type `True`. Knowing that, the type of x should be split accordingly into $\{\text{True}, \neg\text{True}\}$.

This mechanism of backward propagation of splits is initiated in the $[\text{BINDKEEP}_2]$ rule with the two premises $\Gamma \vdash_{\mathcal{E}} (a : \neg(\mathbf{u} \wedge \Gamma'(x))) \Rightarrow \mathbb{F}_1$ and $\Gamma \vdash_{\mathcal{E}} (a : \neg(\mathbf{u} \setminus \Gamma'(x))) \Rightarrow \mathbb{F}_2$. This auxiliary judgment $\Gamma \vdash_{\mathcal{E}} (a : \mathbf{u}) \Rightarrow \mathbb{F}$, defined in Appendix H.3, can be read as follows: *refining the current environment Γ with one of the $\Gamma' \in \mathbb{F}$ ensures that the atom a will have type \mathbf{u}* . The refinements we obtain are stored in the annotation of the binding, using an annotation `propagate` $(\mathcal{A}, \mathbb{F}, \mathcal{S}, \mathcal{S}')$. This annotation is handled by two other rules (omitted here) whose role is to propagate these refinements one after the other using successive `Split` $(\Gamma', \mathcal{K}_1, \mathcal{K}_2)$ results (with $\Gamma' \in \mathbb{F}$), before finally restoring a `keep` $(\mathcal{A}, \mathcal{S}, \mathcal{S}')$ annotation.

$$\begin{array}{c}
\begin{array}{c}
[\text{INTEREMPTY}] \frac{}{\Gamma \vdash_{\mathcal{R}} \langle \eta \mid \wedge(\emptyset, \emptyset) \rangle \Rightarrow \text{Fail}} \quad [\text{INTEROK}] \frac{}{\Gamma \vdash_{\mathcal{R}} \langle \eta \mid \wedge(\emptyset, S) \rangle \Rightarrow \text{Ok}(\wedge(\emptyset, S))} \\
[\text{INTER}_1] \frac{\Gamma \vdash_{\mathcal{R}}^* \langle \eta \mid \mathcal{H} \rangle \Rightarrow \text{Ok}(\mathcal{H}') \quad \Gamma \vdash_{\mathcal{R}} \langle \eta \mid \wedge(S, \{\mathcal{H}'\} \cup S') \rangle \Rightarrow \mathbb{R}}{\Gamma \vdash_{\mathcal{R}} \langle \eta \mid \wedge(\{\mathcal{H}\} \cup S, S') \rangle \Rightarrow \mathbb{R}} \\
[\text{INTER}_2] \frac{\Gamma \vdash_{\mathcal{R}}^* \langle \eta \mid \mathcal{H} \rangle \Rightarrow \text{Fail} \quad \Gamma \vdash_{\mathcal{R}} \langle \eta \mid \wedge(S, S') \rangle \Rightarrow \mathbb{R}}{\Gamma \vdash_{\mathcal{R}} \langle \eta \mid \wedge(\{\mathcal{H}\} \cup S, S') \rangle \Rightarrow \mathbb{R}} \\
[\text{INTER}_3] \frac{\Gamma \vdash_{\mathcal{R}}^* \langle \eta \mid \mathcal{H} \rangle \Rightarrow \mathbb{R}}{\Gamma \vdash_{\mathcal{R}} \langle \eta \mid \wedge(\{\mathcal{H}\} \cup S, S') \rangle \Rightarrow \text{map}(X \mapsto (\wedge(\{X\} \cup S, S')), \mathbb{R})}
\end{array}
\end{array}$$

Intersection annotations are introduced by the $\vdash_{\mathcal{R}}^*$ judgments defined below. In an intersection annotation $\wedge(S, S')$, the annotations in S' are fully processed (i.e., the associated reconstruction returned `Ok()`), while the annotations in S are not: they still have to be refined one after the other (rule $[\text{INTER}_3]$). If one of them becomes fully processed, it is moved in S' (rule $[\text{INTER}_1]$). Conversely, if one of them fails, it is removed (rule $[\text{INTER}_2]$). The process stops when S is empty: then, the reconstruction fails if S' is empty (rule $[\text{INTEREMPTY}]$), and succeed otherwise (rule $[\text{INTEROK}]$).

Finally, we formalize the rules for the judgments $\vdash_{\mathcal{R}}^*$. As said earlier, the purpose of $\vdash_{\mathcal{R}}^*$ is to repeatedly call $\vdash_{\mathcal{R}}$ judgments so that, in the end, we obtain a terminal result.

$$\begin{aligned} \text{[ITERATE}_1\text{]} & \frac{\Gamma \vdash_{\mathcal{R}} \langle \eta \mid \mathcal{H} \rangle \Rightarrow \text{Split}(\Gamma', \mathcal{H}_1, \mathcal{H}_2) \quad \Gamma \vdash_{\mathcal{R}}^* \langle \eta \mid \mathcal{H}_1 \rangle \Rightarrow \mathbb{R}}{\Gamma \vdash_{\mathcal{R}}^* \langle \eta \mid \mathcal{H} \rangle \Rightarrow \mathbb{R}} \quad \Gamma' = \emptyset \\ \text{[ITERATE}_2\text{]} & \frac{\Gamma \vdash_{\mathcal{R}} \langle \eta \mid \mathcal{H} \rangle \Rightarrow \text{Subst}(\{\psi_i\}_{i \in I}, \mathcal{H}_1, \mathcal{H}_2) \quad \Gamma \vdash_{\mathcal{R}}^* \langle \eta \mid \bigwedge(\{\mathcal{H}_1 \psi_i\}_{i \in I} \cup \{\mathcal{H}_2\}, \emptyset) \rangle \Rightarrow \mathbb{R}}{\Gamma \vdash_{\mathcal{R}}^* \langle \eta \mid \mathcal{H} \rangle \Rightarrow \mathbb{R}} \quad \forall i \in I. \psi_i \# \Gamma \end{aligned}$$

The iteration continues as long as it yields non-terminal results that are immediately usable, that is, either they return a trivial split (i.e., $\Gamma' = \emptyset$) as in rule [ITERATE₁], or they return substitutions that do not affect the current environment (i.e., $\psi_i \# \Gamma$) as in rule [ITERATE₂]. For the latter rule, the iteration may need to introduce an intersection annotation (useless when I is empty) in order to explore all the cases of a $\text{Subst}(\{\psi_i\}_{i \in I}, \mathcal{H}_1, \mathcal{H}_2)$ result (where $\mathcal{H}\psi$ is the intermediate annotation \mathcal{H} in which the substitution ψ has been applied recursively to every type in it). An important special case of the [ITERATE₂] rule is when $I = \emptyset$: in that case the iteration continues by trying to type η with the default annotation \mathcal{H}_2 and the current environment Γ . For instance, this special case triggers a [CASEVAR₁] after a [CASETHEN] and a [CASEVAR₂] after a [CASEELSE].

If the result is already terminal or if it is not immediately usable, then it is directly returned:

$$\text{[STOP]} \frac{\Gamma \vdash_{\mathcal{R}} \langle \eta \mid \mathcal{H} \rangle \Rightarrow \mathbb{R}}{\Gamma \vdash_{\mathcal{R}}^* \langle \eta \mid \mathcal{H} \rangle \Rightarrow \mathbb{R}}$$

In particular, if $\mathbb{R} = \text{Split}(\Gamma', \mathcal{H}_1, \mathcal{H}_2)$ where $\Gamma' \neq \emptyset$ (i.e., [ITERATE₁] does not apply), then [STOP] backtracks until Γ' becomes empty; likewise if $\mathbb{R} = \text{Subst}(\{\psi_i\}_{i \in I}, \mathcal{H}_1, \mathcal{H}_2)$ and $\Gamma \psi_i \neq \Gamma$ for some i (i.e. [ITERATE₂] does not apply), then [STOP] backtracks until it exits the scope of the binders of the variables that make the side condition of [ITERATE₂] fail.

4.3 Auxiliary Reconstruction Algorithm

The auxiliary reconstruction algorithm defined in this section converts an intermediate annotation of the main reconstruction system into an annotation for the algorithmic type system. For that, it needs to retrieve the polymorphic substitutions Σ needed to type the atoms.

Formally, the algorithm takes as input an environment Γ , an atom or canonical form η , and an intermediate annotation \mathcal{H} , and produces an annotation \mathfrak{h} for the algorithmic type system. It is presented as a deduction rule system for judgments of the form $\Gamma \vdash_{\mathcal{P}} \langle \eta \mid \mathcal{H} \rangle \Rightarrow \mathfrak{h}$. Some rules have been omitted for concision (they can be found in Appendix H.2): for instance, the rules for constants and axioms are omitted since straightforward, as they just transform an intermediate annotation typ into an annotation \emptyset for the algorithmic type system; likewise, the rules for λ -abstractions and intersections are straightforward and have been omitted, since they just proceed recursively on their children annotations. The most important rule for this system is the one for applications:

$$\text{[APP]} \frac{\begin{array}{c} t_1 = \Gamma(x_1) \quad t_2 = \Gamma(x_2) \\ \rho_1 = \text{refresh}(t_1) \quad \rho_2 = \text{refresh}(t_2) \quad \Sigma = \text{tally}(\{t_1 \rho_1 \dot{\leq} t_2 \rho_2 \rightarrow \alpha\}) \end{array}}{\Gamma \vdash_{\mathcal{P}} \langle x_1 x_2 \mid \text{typ} \rangle \Rightarrow @(\{\sigma \circ \rho_1 \mid \sigma \in \Sigma\}, \{\sigma \circ \rho_2 \mid \sigma \in \Sigma\})} \quad \begin{array}{c} \Sigma \neq \emptyset \\ \alpha \in \mathcal{V}_p \text{ fresh} \end{array}$$

where $\text{refresh}(t)$ returns a renaming from $\text{vars}(t) \cap \mathcal{V}_p$ to fresh polymorphic variables.

For applications, an annotation of the form $@(\Sigma_1, \Sigma_2)$ must be produced. In order to find some instantiations Σ_1 and Σ_2 (for x_1 and x_2 respectively) that make the application typable, the [APP] rule solves the tallying instance $\text{tally}(\{t_1 \rho_1 \dot{\leq} t_2 \rho_2 \rightarrow \alpha\})$. The purpose of ρ_1 and ρ_2 is to decorrelate type variables in $\Gamma(x_1)$ and in $\Gamma(x_2)$. For instance, assume we want to reconstruct the instantiations

for the atom “ x ” with $\Gamma(x) = \beta \rightarrow \beta$. The tallying instance $\text{tally}(\{\beta \rightarrow \beta \dot{\leq} (\beta \rightarrow \beta) \rightarrow \alpha\})$ yields only a very specific, uninteresting solution (i.e., $\alpha = \beta = \mu X.X \rightarrow X$)¹⁰ because of the use of the same type variable β on both sides of $\dot{\leq}$. But each occurrence of x has a polymorphic type that can be instantiated independently. Thus, we remove this useless and constraining dependency by refreshing the generic type variables yielding $\text{tally}(\{\beta' \rightarrow \beta' \dot{\leq} (\beta \rightarrow \beta) \rightarrow \alpha\})$ which has interesting solutions, in particular $\{\beta' \rightsquigarrow \beta \rightarrow \beta; \alpha \rightsquigarrow \beta \rightarrow \beta\}$. The side-condition $\Sigma \neq \emptyset$ ensures that the tallying instance has at least one solution (otherwise the annotation produced would be invalid). The rules for projections, pairs, and type-cases are similar and, thus, omitted.

$$[\text{BINDKEEP}] \frac{\Gamma \vdash_{\mathcal{P}} \langle a \mid \mathcal{A} \rangle \Rightarrow \emptyset \quad \Gamma \vdash_{\mathcal{A}} [a \mid \emptyset] : s \quad (\forall i \in I) \Gamma, x : s \wedge \mathbf{u}_i \vdash_{\mathcal{P}} \langle \kappa \mid \mathcal{K}_i \rangle \Rightarrow \mathbb{k}_i}{\Gamma \vdash_{\mathcal{P}} \langle \text{bind } x = a \text{ in } \kappa \mid \text{keep } (\mathcal{A}, \emptyset, \{\langle \mathbf{u}_i, \mathcal{K}_i \rangle\}_{i \in I}) \rangle \Rightarrow \text{keep } (\emptyset, \{\langle \mathbf{u}_i, \mathbb{k}_i \rangle\}_{i \in I})} (*)$$

(where $(*)$ is $\bigvee_{i \in I} \mathbf{u}_i \simeq \mathbb{1}$). The rule [BINDKEEP] takes as input an intermediate annotation $\text{keep } (\mathcal{A}, \mathcal{S}, \mathcal{S}')$, with $\mathcal{S} = \emptyset$, since all branches must have been fully explored by the main reconstruction algorithm. The rule recursively transforms the intermediate annotation \mathcal{A} for the definition a into an annotation \emptyset for the algorithmic type system, and uses it to type a . It can then update the environment and proceed recursively on the body κ , for each branch in \mathcal{S}' .

4.4 Properties of the Reconstruction Algorithm

As recalled at the beginning of the section, reconstruction is sound, terminating, but incomplete.

THEOREM 4.2 (SOUNDNESS). *If $\Gamma \vdash_{\mathcal{P}} \langle \kappa \mid \mathcal{K} \rangle \Rightarrow \mathbb{k}$, then $\exists t. \Gamma \vdash_{\mathcal{R}} [\kappa \mid \mathbb{k}] : t$.*

THEOREM 4.3 (TERMINATION). *The deduction rules $\vdash_{\mathcal{R}}^*$ and $\vdash_{\mathcal{R}}$ define a terminating algorithm: it can either fail (if no rule applies at some point) or return a result \mathbb{R} .*

The incompleteness of the reconstruction algorithm is inherent to our system and derives from the lack of principal typing. A simple example is the curried function `map` defined in the third row of Table 1 in the next section. Our reconstruction deduces for it the type $(\alpha \rightarrow \beta) \rightarrow [\alpha^*] \rightarrow [\beta^*]$ (actually, a slightly more precise type), where $[\alpha^*]$ is the type of the lists of elements of type α . This states that an application of `map` yields a function that maps lists of α 's into lists of β 's. But for every natural number n , the declarative system can also deduce that the result maps lists of α 's of length n into lists of β 's of the same length n . Our algorithm can *check* each of these types, but none of them can be deduced from the type reconstructed by the algorithm. And since we do not have dependent types or infinite intersections, then the declarative system cannot have a principal type expressing all these different derivations. In other terms, incompleteness stems from the fact that the declarative system can use all the infinitely many decompositions of unions in the union elimination rule, and the infinitely many decompositions of the domain of a function when reconstructing its type as an intersection of arrows. The algorithmic counterpart of this, is that there are infinitely many annotations that the algorithmic system can use to type these expressions and that these infinite choices cannot be summarized by a notion of principal annotation: the reconstruction chooses one particular annotation, and therefore it will miss some solutions.

There is a second source of incompleteness for reconstruction, which is not inherent to the system, but a design choice, instead: the fact that reconstruction does not perform the so-called “expansion” of intersection types. This is shown by the rule [APP] in Section 4.3, where tally is applied without expanding the types in the constraint (e.g., if $\text{tally}(\{t_1\rho_1 \dot{\leq} t_2\rho_2 \rightarrow \alpha\})$ fails we can expand the type of the function and try $\text{tally}(\{t_1\rho_1 \wedge t_1\rho_3 \dot{\leq} t_2\rho_2 \rightarrow \alpha\})$, and so on and so forth by alternating expansions on the function and on the argument types: see [Castagna et al.

¹⁰The solution is not interesting since it is the one that allows any simply typed system to type all pure lambda terms. We do not need this recursive type to type, say, the application of the polymorphic identity function to itself.

Table 1. Types inferred by the implementation (times are in ms)

	Code	Inferred type	Time
1	<pre> type Falsy = False "" 0 type Truthy = ~Falsy let toBoolean x = if x is Truthy then true else false let lOr (x,y) = if toBoolean x then x else y let id x = lOr (x,x) </pre>	$(\text{Falsy} \rightarrow \text{False}) \wedge (\text{Truthy} \rightarrow \text{True})$ $(((\alpha \wedge \text{Truthy}) \times \mathbb{1}) \rightarrow \alpha \wedge \text{Truthy}) \wedge$ $((\text{Falsy} \times \beta) \rightarrow \beta)$ $\alpha \rightarrow \alpha$	3.42 13.31 8.46
2	<pre> let fixpoint = fun f -> let delta = fun x -> f (fun v -> (x x v)) in delta delta </pre>	$((\beta \rightarrow \alpha) \rightarrow (\beta \rightarrow \alpha) \wedge \gamma) \rightarrow (\beta \rightarrow \alpha) \wedge \gamma$	15.37
3	<pre> let map_stub map f lst = if lst is Nil then nil else (f (fst lst), map f (snd lst)) let map = fixpoint map_stub </pre>	... $((\alpha \rightarrow \beta) \rightarrow ([\alpha^*] \rightarrow [\beta^*])) \wedge (\mathbb{1} \rightarrow [] \rightarrow [])$	33.03 84.75
4	<pre> let filter_stub filter (f: ('a->Any) & ('b -> ~True)) (l: [('a 'b)*]) = if l is Nil then nil else if f(fst(l)) is True then (fst(l), filter f (snd(l))) else filter f (snd(l)) let filter = fixpoint filter_stub </pre>	... $((\alpha \rightarrow \mathbb{1}) \wedge (\beta \rightarrow \neg \text{True})) \rightarrow [(\alpha \vee \beta)^*] \rightarrow [(\alpha \wedge \beta)^*]$	21.19 13.83
5	<pre> let rec flatten x = match x with [] -> [] h::t -> concat (flatten h) (flatten t) _ -> [x] </pre>	$(\text{Tree} \rightarrow [(\alpha \setminus [\mathbb{1}^*])^*]) \wedge (\beta \setminus [\mathbb{1}^*] \rightarrow [\beta \setminus [\mathbb{1}^*]])$ where $\text{Tree} = [\text{Tree}^*] \vee (\alpha \setminus [\mathbb{1}^+])$	374.41

2015, Section 3.2.3] for more details). The consequence of this is that if you take the definition of the function `filter` given in row 4 of Table 1, and you remove all type annotations, then the type reconstructed by the algorithm for it is less precise than the one specified by the annotations, which could have been reconstructed if the algorithm had instead expanded the type of the parameter `f`.

Despite incompleteness, the declarative rules of Figure 2 form a reliable guide to which programs are accepted, provided we bear in mind that the algorithm approximates data structures according to the tests performed on them. So, typically, the type reconstructed for a function on lists, will probably differentiate the cases for empty and not-empty lists, but not for, say, lists of size 42, unless the function contains an explicit test for it. This (and to a lesser extent, expansion) is essentially the main difference with the declarative system, which has the liberty to deduce the type for the case of lists of size 42, even if this property is not tested in the body of the function. In that case, the programmer can still use an explicit type annotation to *check* that the specific type works.

5 IMPLEMENTATION

We have implemented the reconstruction algorithm presented in Section 4, using the CDuce [CDuce] API for the subtyping and the tallying algorithms. The prototype is 4500 lines of OCaml code and features several extensions such as optional type annotations, pattern matching (cf. Appendix A), records, and a more user-friendly syntax. It implements some optimizations, briefly discussed at the end of this section, for instance memoization and a mechanism to avoid typing redundant branches when inferring the domains of λ -abstractions. We give in Table 1 the code of several functions, using a syntax similar to OCaml, where uppercase identifiers (e.g., `True`, `Truthy`) denote types and lowercase identifiers denote variables or constants. For each function we report its inferred type and the time used to infer it. To enhance readability we manually curated the types which, thus, may be syntactically different from (but are semantically equivalent to) the types printed by the prototype. The experiments were performed on an Intel Core i9-10900KF 3.70GHz CPU. The code was compiled natively using OCaml 4.14.1. All these examples (and more) can be tested on the

web-based interactive prototype hosted at <https://www.cduce.org/dynlang>. The web version is compiled to JavaScript using js_of_ocaml [Ocsigen], and is about 8 times slower than the native version.

Code 1 features the examples used in the introduction.

Code 2 implements Curry’s fix-point combinator in a call-by-value setting. Though it is traditionally given the type $((\beta \rightarrow \alpha) \rightarrow (\beta \rightarrow \alpha)) \rightarrow (\beta \rightarrow \alpha)$, our prototype infers a slightly more precise type by intersecting the co-domain of the argument with γ .

Code 3 shows how to use the fix-point combinator to type recursive functions. The `map_stub` function implements a step of the traditional `map` function. The type inferred for this function has been omitted for simplicity. Then, `map` is obtained by applying the fixed-point combinator to `map_stub`. Note that $[\alpha^*]$ denotes a list of elements of type α , and $[]$ denotes an empty list. The branch $\mathbb{1} \rightarrow [] \rightarrow []$ may be surprising, but it is correct since the `map` function does not use its first argument if the second argument is an empty list.

Code 4 shows how type annotations (cf. Appendix A.2) can be used to infer more precise types: when the `filter` function is applied to a characteristic function for the set $\alpha \vee \beta$ whose type precises that the elements in β do not satisfy the predicate, then the inferred type has these elements removed from the type of the result.

The grammar for expressions in Figure 1 does not include recursive functions, since *from a theoretical viewpoint* they are useless: Milner [1978, page 356] justifies the addition of a “fix $x.e$ ” expression by the fact that his system cannot type Curry’s fixpoint combinator, but, as explained in Section 1.1 and shown by Code 2 above, our system can. However, *from a practical viewpoint*, the use of `let rec` definitions instead of fixed-points combinators may dramatically improve the speed of reconstruction, which is why the previous definitions of `map` and `filter` with a fixed-point combinator must be considered just as stress tests for our reconstruction algorithm. For recursive functions we implemented classic `let rec` definitions, for which the reconstruction takes the arity of the function into account. Code 5 shows the use of `let rec` and of pattern matching (cf. Appendix A.3) and is an example of the improvement brought by `let rec` definitions: reconstruction for the same definition but with a fixpoint combinator is four times slower. The code defines the deep `flatten` function that transforms arbitrary nested lists into the list of their elements (where `concat` is a function of type $[\alpha^*] \rightarrow [\beta^*] \rightarrow [(\alpha^*)(\beta^*)]$, the result being the type of lists starting with α elements and ending with β ones). Greenberg [2019] considers this function to be the ultimate test for any type system: as he explains, this simple polymorphic function defies all type systems since of all existing languages, none can reconstruct a type for it and only a couple of languages can check its explicitly typed version: CDuce and Haskell (the latter by resorting to complex metaprogramming constructions). Our system reconstructs a precise type for `flatten` as shown by the first arrow in its intersection type, which states that `flatten` is a function that takes a tree (i.e., either a list of elements that are trees, or a value different from a list) and returns the list of elements of the tree that are not lists; the other arrow of the intersection states that when `flatten` is applied to an element different from a list, then it returns the list containing only that element.

Our prototype focuses on proximity with the inference system for reconstruction, rather than on performance: we used it mainly to explore and test our system, which is why it is implemented in a purely functional style with persistent data structures (so as to simulate the reconstruction inference rules). Nonetheless, a few optimizations were implemented in order to mitigate the cost of backtracking and branching. One source of inefficiency comes from the intersection nodes that are generated when a destructor is reconstructed. This generation can lead to an explosion of the number of branches to explore, even though many of these branches are redundant. In the prototype, this is mitigated by recording, for each λ -abstraction, the domains already explored for it, and by trimming branches that do not explore new combinations of domains.

Another source of inefficiency comes from the type decompositions performed after each binding. Although these type decompositions are usually small (e.g., the type of a binding is seldom split in more than two parts), it becomes an issue when typing large expressions with multiple type-cases. For instance, a preliminary and unoptimized implementation of the reconstruction algorithm took about 40 seconds to type the `bal(ance)` function used in the module `Map` of the OCaml standard library, that contains 6 different pattern matches and 4 type-cases [OCaml 2023]. Adding a simple memoization mechanism that prevents the reconstruction from retyping an atom several times for equivalent contexts, decreased the inference time down to 4 seconds.

While these simple optimizations significantly improve performance, they are still far from what would be considered acceptable for real applications. To be used in mainstream languages, the type system will have to be adapted and restricted so as to ensure better and uniform performance. To this purpose, we believe that some more language-oriented optimization techniques could be of help. An example is what the development team of Luau [Luau] did on the occasion of its recent switch to semantic subtyping [Jeffrey 2022]. The developers did this switch by implementing a two-phase approach: first, a sound syntactic system, fast but imprecise, is used to try to prove subtyping, and only if it fails, the computationally expensive semantic subtyping inference is used. We think not only that such a staged approach could be applied in our case, but also that the partial results of the first phase could be used to improve the performance of the later phases, as in the case of the `let rec`, where knowing the arity of the defined function improves the performance of the reconstruction. This could be further coupled with slicing, meaning that our type reconstruction could be applied to very delimited regions that would bound the possibility of backtracking. These techniques are language-dependent, and quite different from the algorithmic aspects developed here, though they will completely rely on it. We plan to explore them in future work.

6 RELATED WORK

This work can be seen as a polymorphic extension of [Castagna et al. 2022b] from which it borrows some key notions, such as (i) the combination of the union elimination rule (from [Barbanera et al. 1995]) with three rules for type-cases, in order to capture the essence of occurrence typing ([Tobin-Hochstadt and Felleisen 2008]), (ii) the use of MSC forms to drive the application of the union elimination rule, and (iii) the use of annotations in the algorithmic type system. However, the introduction of polymorphic types greatly modifies the meta-theory. Besides its influence on the union elimination rule, the interplay between intersection, union elimination and instantiation suggests a different style of type annotations, to be amenable to type inference. We use external annotations while [Castagna et al. 2022b] annotates terms. Further, the presence of type variables imposes to use tallying in an inference algorithm inspired by \mathcal{W} by Damas and Milner [1982] and from [Castagna et al. 2015], where tallying was first introduced to type polymorphic applications. This yields a clear improvement over [Castagna et al. 2022b] which is unable to infer higher-order types for function arguments, while our algorithm is able to do so even for recursive functions.

The use of trees to annotate calculi with full-fledged intersection types is common. In the presence of explicitly-typed overloaded functions, one must be able to precisely describe how the types of nested λ -abstractions relate to the various “branches” of the outermost function. The work most similar to ours is [Liquori and Ronchi Della Rocca 2007], since the deductions are performed on pairs of marked term and proof term. A marked term is an untyped term where variables are marked with integers and a proof term is a tree that encodes the structure of the typing derivation and relates marks to types. Other approaches, such as [Bono et al. 2008; Ronchi Della Rocca 2002; Wells et al. 2002], duplicate the term typed with an intersection, such that each copy corresponds exactly to one member of the intersection. Lastly, the work of [Wells and Haack 2002] does not duplicate terms but rather decorate λ -abstractions with a richer concept of *branching shape* which

essentially allows one to give names to the various branches of an overloaded function and to use these names in the annotations of nested λ -abstraction. Note that none of these works features type reconstruction, which was our main motivation to eschew annotations within terms, since the backtracking nature of our reconstruction would imply rewriting terms over and over.

Inference for ML systems with subtyping, unions, and intersections has been studied in MLsub [Dolan and Mycroft 2017] and extended with richer types and a limited form of negation in MLstruct [Parreaux and Chau 2022]. Both works trade expressivity for principality. They define a lattice of types and an algebraic subtyping relation that ensures principality, but forbids the intersection of arrow types. This precludes them from expressing overloaded functions, but allows them to define a principal polymorphic type inference with unions and intersections. We justify our choice of set-theoretic types, with no type principality and a complex inference, by our aim to type dynamic languages, such as Erlang or JavaScript, where overloading plays an important role. We favour the expressivity necessary to type many idioms of these languages, and rely on user-defined annotations when necessary to compensate for the incompleteness of type inference. Lastly, both works implement some form of type simplifications (e.g., Dolan and Mycroft [2017] use automata techniques to simplify types), a problem of practical importance that we did not tackle, yet.

Angelo and Florido [2022] provide a principal type inference for a type system with rank-2 intersection types. In their work, overloaded behaviors are expressible using intersection types, but they are limited by the rank-2 restriction. Union types are not supported, nor are equi-recursive types (actually, it does not feature a general notion of subtyping between two arbitrary types). Their inference does not require backtracking: it generates a set of constraints that are then solved using a *set unification algorithm*. This approach for inference has some similarities with the one by Castagna et al. [2016] improved and further developed by Petrucciani [2019] in a context with set-theoretic types, where the *set unification algorithm* is replaced by *tallying* in the presence of subtyping. However, while [Petrucciani 2019] does support intersection types with no ranking limitation, it is not able to infer intersection types for overloaded functions. Our work aims to improve this aspect, as well as providing a more precise typing of type-cases (occurrence typing).

Work by Oliveira et al. [2016] and Rioux et al. [2023] study disjoint intersection and union types. They allow expressing overloaded behaviors by a general deterministic merge operator. In our work, we do not have a general merge operator: overloaded behaviors only emerge through the use of type-case expressions (or the application of an overloaded function). Our work can be extended with pattern-matching, in which case the first matching branch is selected. This is a different approach than the one used with disjoint intersection types, where branches are disjoint and have no priority and where ambiguous programs are rejected using a notion of mergeability and distinguishability, allowing to define a general merge operator and to support nested composition, which may be useful in some contexts such as compositional programming [Zhang et al. 2021].

Jim [2000] presents a polar type system which features intersections and parametric polymorphism. In Jim's type system, quantifiers may appear only in positive positions in types, while intersections may only appear in negative positions. This yields a system that is more expressive than rank-2 intersection types, and therefore more expressive than ML. Furthermore, the system features principal types, and a decidable type inference. Some aspects of this work are similar to ours, in particular the use of MGS, an algorithm to compute the most general solution of a (syntactic) sub-typing problem, that plays the same role as our tallying algorithm. Despite these similarities, the approaches differ in the kind of programs they handle: in [Jim 2000], intersections are only deduced by applying higher-order function parameters to arguments of distinct types within the body of a function, while in our approach, they can also be caused by a type-case.

Finally, set-theoretic types are starting to be integrated into *real-world languages*, for instance by Schimpf et al. [2023] for Erlang, by Jeffrey [2022] for Lua, and by Castagna et al. [2023a] for Elixir.

We believe that, in the future, our work could be used in these systems in order to benefit from a more precise typing of type-cases and pattern matching, as well as by providing an optional type inference that can be used in conjunction with explicit type annotations.

7 CONCLUSION

This work aims at providing a formal and expressive type system for dynamic languages, where type-cases can be used to give functions an overloaded behavior. It features a type inference that mixes both parametric polymorphism (for modularity) and intersection polymorphism (to capture overloaded behaviors). In that sense, our work is more than a simple study on typability: as a matter of fact, monomorphic intersection and union types are sufficient to type a closed program where all function applications are known (cf., Section 1.2), but this would be bad from a language design point of view, and it is the reason why people program using ML-style programming languages rather than intersection based ones. Separate compilation and modular definitions are requirements of any reasonable programming language. The essence of this work is thus to challenge the limits of how much precision one can obtain (through intersection types)—ideally precise enough to type idioms of dynamic languages—while preserving modularity (thanks to let-polymorphism).

While we believe our work to be an important step towards a better static typing of dynamic languages, several key features are still missing. First, the presence of side effects may invalidate our approach: if the $[\vee]$ rule in Figure 2 is applied to two different occurrences of an expression e' that is not pure, then the rule may type an expression that yields a run-time type error. This can be seen on the algorithmic system, where the transformation into an MSC-form binds the two occurrences of e' to the same variable, thus wrongly assuming that they both yield the same result. Strictly speaking, our algorithmic approach does not require expressions to be pure; it just needs that when two occurrences of an expression may produce two distinct values that may change the result of a dynamic test, then these two occurrences must be bound by two different binds. Having only pure expressions is a straightforward way to satisfy this property. Having each subexpression bound to a distinct variable (i.e., no sharing, that is, a less precise system, in which the union rule is never used) is a way to retain safety in the presence of side-effects. But between these two extrema, there is a whole palette of less coarse solutions that make it possible to apply our approach in the presence of side-effects, and that we plan to study in future work. This poses two main challenges: (i) how to separate problematic expressions from non-problematic ones (e.g., a `gen_id: Unit → Int` function performs side-effects, but if its result is tested only against `Int`, then it is sound to have all occurrences of `gen_id` bound by the same bind during typing) which, in terms of the type system, corresponds to characterize a class of subexpressions e' that can be safely used in rule $[\vee]$; and (ii) how to do so *before* our type inference, at a point when type information is not available, yet.

Second, while the performance of our prototype is reasonable, it can certainly be improved by using more sophisticated implementations techniques and heuristics on the lines we outlined at the end of Section 5.

Third, the interactions between code that is exported and code that is local must be better studied and understood: using intersection for local polymorphic functions and generalization for global ones, may not always be entirely satisfactory since the types of the global functions may be “polluted” by the types of the local applications, yielding less a precise reconstruction for the former. One solution can be to hoist the definition of polymorphic functions at toplevel whenever possible.

Lastly, an important future work is the support of row-polymorphism: while records can be easily added to the present work, the precise typing of functions operating on records requires row-polymorphism. This is especially important for dynamic languages where records are seamlessly used to encode both objects and dictionaries. A first step in that direction may be to integrate the work by Castagna [2023b], which unifies dictionaries and records.

DATA AVAILABILITY STATEMENT

- All the auxiliary definitions, proofs and extensions that we omitted from the main text are available in the appendix of the extended version [Castagna et al. 2024].
- The reference implementation of the type-reconstruction system is archived on Zenodo [Castagna et al. 2023b]. An online version of the prototype is available at:

<https://www.cduce.org/dynlang>

ACKNOWLEDGMENTS

We warmly thank the POPL reviewers: their careful reading and suggestions allowed us to improve the presentation significantly. A special thank the reviewers of the POPL artifact evaluation for their detailed and insightful reviews.

This work was partially supported by the *Chaire Langages Dynamiques pour les Données* of the *Fondation Université Paris-Saclay*, by the *SECUREVAL* ANR project n. ANR-22-PECY-0005 and by a CIFRE PhD. grant with Remote Technology.

REFERENCES

- Pedro Ângelo and Mário Florido. 2022. Type Inference For Rank-2 Intersection Types Using Set Unification. In *Theoretical Aspects of Computing – ICTAC 2022: 19th International Colloquium, Tbilisi, Georgia, September 27–29, 2022, Proceedings* (Tbilisi, Georgia). Springer-Verlag, Berlin, Heidelberg, 462–480. https://doi.org/10.1007/978-3-031-17715-6_29
- Franco Barbanera, Mariangiola Dezani-Ciancaglini, and Ugo de'Liguoro. 1995. Intersection and Union Types. *Inf. Comput.* 119, 2 (June 1995), 202–230. <https://doi.org/10.1006/inco.1995.1086>
- Viviana Bono, Betti Venneri, and Lorenzo Bettini. 2008. A typed lambda calculus with intersection types. *Theor. Comput. Sci.* 398, 1-3 (2008), 95–113. <https://doi.org/10.1016/j.tcs.2008.01.046>
- Giuseppe Castagna. 2023a. Programming with union, intersection, and negation types. In *The French School of Programming*, Bertrand Meyer (Ed.). Springer. ISBN 978-3-031-34517-3. Preprint at [arXiv:2111.03354](https://arxiv.org/abs/2111.03354).
- Giuseppe Castagna. 2023b. Typing Records, Maps, and Structs. *Proc. ACM Program. Lang.* 7, ICFP, Article 196 (Sept. 2023), 45 pages. <https://doi.org/10.1145/3607838>
- Giuseppe Castagna, Guillaume Duboc, and José Valim. 2023a. The Design Principles of the Elixir Type System. *The Art, Science, and Engineering of Programming* 8, 2 (2023). <https://doi.org/10.22152/programming-journal.org/2024/8/4> Preprint in ArXiv: [arXiv:2306.06391](https://arxiv.org/abs/2306.06391).
- Giuseppe Castagna, Victor Lanvin, Mickaël Laurent, and Kim Nguyen. 2022a. Revisiting Occurrence Typing. *Science of Computer Programming* 217 (mar 2022), 102781. <https://doi.org/10.1016/j.scico.2022.102781> [arXiv:1907.05590](https://arxiv.org/abs/1907.05590)
- Giuseppe Castagna, Mickaël Laurent, and Kim Nguyen. 2023b. *Prototype: Polymorphic Type Inference for Dynamic Languages*. <https://doi.org/10.5281/zenodo.10155221>
- Giuseppe Castagna, Mickaël Laurent, and Kim Nguyen. 2024. Polymorphic Type Inference for Dynamic Languages (Extended version). <https://doi.org/10.1145/3632882>
- Giuseppe Castagna, Mickaël Laurent, Kim Nguyen, and Matthew Lutze. 2022b. On Type-Cases, Union Elimination, and Occurrence Typing. *Proc. ACM Program. Lang.* 6, POPL, Article 13 (jan 2022), 31 pages. <https://doi.org/10.1145/3498674>
- Giuseppe Castagna, Kim Nguyen, Zhiwu Xu, and Pietro Abate. 2015. Polymorphic functions with set-theoretic types. Part 2: local type inference and type reconstruction. In *Proceedings of the 42nd Annual ACM SIGPLAN Symposium on Principles of Programming Languages (POPL '15)*. 289–302. <https://doi.org/10.1145/2676726.2676991>
- Giuseppe Castagna, Tommaso Petrucciani, and Kim Nguyen. 2016. Set-Theoretic Types for Polymorphic Variants. In *ICFP '16, 21st ACM SIGPLAN International Conference on Functional Programming*. 378–391. <https://doi.org/10.1145/2951913.2951928>
- Giuseppe Castagna and Zhiwu Xu. 2011. Set-theoretic Foundation of Parametric Polymorphism and Subtyping. In *ICFP '11: 16th ACM-SIGPLAN International Conference on Functional Programming*. 94–106. <https://doi.org/10.1145/2034773.2034788>
- CDuce. *The CDuce Compiler*. CDuce <https://www.cduce.org>
- Mario Coppo, Mariangiola Dezani-Ciancaglini, and Betti Venneri. 1981. Functional Characters of Solvable Terms. *Mathematical Logic Quarterly* 27, 2-6 (1981), 45–58. <https://doi.org/10.1002/malq.19810270205>
- Luis Damas and Robin Milner. 1982. Principal Type-Schemes for Functional Programs. In *Proceedings of the 9th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)* (Albuquerque, New Mexico). Association for Computing Machinery, New York, NY, USA, 207–212. <https://doi.org/10.1145/582153.582176>
- Mariangiola Dezani. 2020. Personal communication.

- Stephen Dolan and Alan Mycroft. 2017. Polymorphism, Subtyping, and Type Inference in MLsub. In *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages (Paris, France) (POPL '17)*, Giuseppe Castagna and Andrew D. Gordon (Eds.), Association for Computing Machinery, New York, NY, USA, 60–72. <https://doi.org/10.1145/3009837.3009882>
- Ecma. 2021. ECMAScript® 2021 Language Specification. <https://262.ecma-international.org/12.0/>
- Facebook. *Flow*. Facebook <https://flow.org/>
- Alain Frisch, Giuseppe Castagna, and Véronique Benzaken. 2002. Semantic Subtyping. In *LICS '02, 17th Annual IEEE Symposium on Logic in Computer Science*. IEEE Computer Society Press, 137–146. <https://doi.org/10.1109/LICS.2002.1029823>
- Alain Frisch, Giuseppe Castagna, and Véronique Benzaken. 2008. Semantic subtyping: dealing set-theoretically with function, union, intersection, and negation types. *Journal of the ACM* 55, 4 (Sept. 2008), 19:1–19:64. <https://doi.org/10.1145/1391289.1391293>
- Nils Gesbert, Pierre Genevès, and Nabil Layaida. 2015. A logical approach to deciding semantic subtyping. *ACM Transactions on Programming Languages and Systems* 38, 1 (2015), 3. <https://doi.org/10.1145/2812805>
- Michael Greenberg. 2019. The Dynamic Practice and Static Theory of Gradual Typing. In *3rd Summit on Advances in Programming Languages (SNAPL 2019) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 136)*. 6:1–6:20. <https://doi.org/10.4230/LIPIcs.SNAPL.2019.6>
- Fritz Henglein. 1993. Type Inference with Polymorphic Recursion. *ACM Trans. Program. Lang. Syst.* 15, 2 (apr 1993), 253–289. <https://doi.org/10.1145/169701.169692>
- J. Roger Hindley. 1969. The Principal Type-Scheme of an Object in Combinatory Logic. *Trans. Amer. Math. Soc.* 146 (1969), 29–60. <https://doi.org/10.2307/1995158>
- Alan Jeffrey. 2022. Semantic Subtyping in Luau. Roblox Technical Blog. <https://blog.roblox.com/2022/11/semantic-subtyping-luau> Accessed on May 6th 2023.
- Trevor Jim. 2000. A Polar Type System. In *ICALP Workshops 2000, Proceedings of the Satellite Workshops of the 27th International Colloquium on Automata, Languages and Programming, Geneva, Switzerland, July 9-15, 2000*, José D. P. Rolim, Andrei Z. Broder, Andrea Corradini, Roberto Gorrieri, Reiko Heckel, Juraj Hromkovic, Ugo Vaccaro, and Joe B. Wells (Eds.). Carleton Scientific, Waterloo, Ontario, Canada, 323–338.
- Assaf J. Kfoury, Jerzy Tiurny, and Pawel Urzyczyn. 1993. Type Reconstruction in the Presence of Polymorphic Recursion. *ACM Trans. Program. Lang. Syst.* 15, 2 (apr 1993), 290–311. <https://doi.org/10.1145/169701.169687>
- Daniel Leivant. 1983. Polymorphic Type Inference. In *Proceedings of the 10th ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages (POPL)* (Austin, Texas). Association for Computing Machinery, New York, NY, USA, 88–98. <https://doi.org/10.1145/567067.567077>
- Luigi Liquori and Simona Ronchi Della Rocca. 2007. Intersection-types à la Church. *Inf. Comput.* 205, 9 (2007), 1371–1386. <https://doi.org/10.1016/j.ic.2007.03.005>
- Luau. *Luau*. <https://luau-lang.org/>
- David MacQueen, Gordon Plotkin, and Ravi Sethi. 1986. An ideal model for recursive polymorphic types. *Information and Control* 71, 1 (1986), 95–130. [https://doi.org/10.1016/S0019-9958\(86\)80019-5](https://doi.org/10.1016/S0019-9958(86)80019-5)
- Per Martin-Löf. 1994. *Analytic and Synthetic Judgements in Type Theory*. Springer Netherlands, Dordrecht, 87–99. https://doi.org/10.1007/978-94-011-0834-8_5
- Microsoft. *TypeScript*. Microsoft <https://www.typescriptlang.org/>
- Robin Milner. 1978. A theory of type polymorphism in programming. *J. Comput. System Sci.* 17, 3 (1978), 348–375. [https://doi.org/10.1016/0022-0000\(78\)90014-4](https://doi.org/10.1016/0022-0000(78)90014-4)
- OCaml. 2023. Standard Library: Map module. Github repository. <https://github.com/ocaml/ocaml/blob/trunk/stdlib/map.ml>
- Ocsigen. *JS of OCaml*. Ocsigen https://ocsigen.org/js_of_ocaml/
- Bruno C. d. S. Oliveira, Zhiyuan Shi, and João Alpuim. 2016. Disjoint Intersection Types. *SIGPLAN Not.* 51, 9 (sep 2016), 364–377. <https://doi.org/10.1145/3022670.2951945>
- Lionel Parreaux and Chun Yin Chau. 2022. MLstruct: Principal Type Inference in a Boolean Algebra of Structural Types. *Proc. ACM Program. Lang.* 6, OOPSLA2, Article 141 (oct 2022), 30 pages. <https://doi.org/10.1145/3563304>
- Tommaso Petrucciani. 2019. *Polymorphic Set-Theoretic Types for Functional Languages*. Ph. D. Dissertation. Joint Ph.D. Thesis, Università di Genova and Université Paris Diderot. <https://tel.archives-ouvertes.fr/tel-02119930>
- Tommaso Petrucciani, Giuseppe Castagna, Davide Ancona, and Elena Zucca. 2018. Semantic Subtyping for Non-Strict Languages. In *TYPES18: 24th International Conference on Types for Proofs and Programs (LIPIcs, Vol. 130)*, Peter Dybjer, José Espírito Santo, and Luís Pinto (Eds.). 4:1–4:24. <https://doi.org/10.4230/LIPIcs.TYPES.2018.4>
- Nick Rioux, Xuejing Huang, Bruno C. d. S. Oliveira, and Steve Zdancewic. 2023. A Bowtie for a Beast: Overloading, Eta Expansion, and Extensible Data Types in F_⋄. *Proc. ACM Program. Lang.* 7, POPL, Article 18 (jan 2023), 29 pages. <https://doi.org/10.1145/3571211>

- John Alan Robinson. 1965. A Machine-Oriented Logic Based on the Resolution Principle. *J. ACM* 12, 1 (jan 1965), 23–41. <https://doi.org/10.1145/321250.321253>
- Simona Ronchi Della Rocca. 2002. Intersection Typed lambda-calculus. *Electr. Notes Theor. Comput. Sci.* 70, 1 (2002), 163–181. [https://doi.org/10.1016/S1571-0661\(04\)80496-1](https://doi.org/10.1016/S1571-0661(04)80496-1)
- Amr Sabry and Matthias Felleisen. 1992. Reasoning about Programs in Continuation-Passing Style.. In *Proceedings of the 1992 ACM Conference on LISP and Functional Programming (LFP '92)*. Association for Computing Machinery, 288–298. <https://doi.org/10.1145/141471.141563>
- Albert Schimpf, Stefan Wehr, and Annette Bieniusa. 2023. Set-Theoretic Types for Erlang. In *Proceedings of the 34th Symposium on Implementation and Application of Functional Languages (Copenhagen, Denmark) (IFL '22)*. Association for Computing Machinery, New York, NY, USA, Article 4, 14 pages. <https://doi.org/10.1145/3587216.3587220>
- Christopher Strachey. 1967. Fundamental concepts in programming languages. Lecture notes for International Summer School in Computer Programming, Copenhagen.
- Sam Tobin-Hochstadt and Matthias Felleisen. 2008. The Design and Implementation of Typed Scheme. In *Proceedings of the 35th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (San Francisco, California, USA) (POPL '08)*. ACM, New York, NY, USA, 395–406. <https://doi.org/10.1145/1328438.1328486>
- Sam Tobin-Hochstadt and Matthias Felleisen. 2010. Logical types for untyped languages. In *Proceedings of the 15th ACM SIGPLAN International Conference on Functional Programming (Baltimore, Maryland, USA) (ICFP '10)*. ACM, New York, NY, USA, 117–128. <https://doi.org/10.1145/1863543.1863561>
- Types 2019. What exactly should we call syntax-directed inference rules? Discussion on the Types mailing list. <http://lists.seas.upenn.edu/pipermail/types-list/2019/002138.html>.
- Joseph Brian Wells, Allyn Dimock, Robert J Muller, and Franklyn Albin Turbak. 2002. A calculus with polymorphic and polyvariant flow types. *J. Funct. Program.* 12, 3 (2002), 183–227. <https://doi.org/10.1017/S0956796801004245>
- Joe B. Wells and Christian Haack. 2002. Branching Types. In *ESOP '02 (LNCS, Vol. 2305)*. Springer, 115–132. https://doi.org/10.1007/3-540-45927-8_9
- Weixin Zhang, Yaozhu Sun, and Bruno C. D. S. Oliveira. 2021. Compositional Programming. *ACM Trans. Program. Lang. Syst.* 43, 3, Article 9 (sep 2021), 61 pages. <https://doi.org/10.1145/3460228>