A Type-Directed Operational Semantics For a Calculus with a Merge Operator

Xuejing Huang

The University of Hong Kong, China
xjhuang@cs.hku.hk

Bruno C. d. S. Oliveira

The University of Hong Kong, China
bruno@cs.hku.hk

──── **Abstract** ────

Calculi with disjoint intersection types and a merge operator provide general mechanisms that can subsume various other features. Such calculi can also encode highly dynamic forms of object composition, which capture common programming patterns in dynamically typed languages (such as JavaScript) in a fully statically typed manner. Unfortunately, unlike many other foundational calculi (such as *System F*, *System F$_{<:}$* or *Featherweight Java*), recent calculi with the merge operator lack a (direct) operational semantics with standard and expected properties such as *determinism* and *subject-reduction*. Furthermore the metatheory for such calculi can only account for *terminating programs*, which is a significant restriction in practice.

This paper proposes a *type-directed operational semantics* (TDOS) for $\lambda_i^{:}$: a calculus with *intersection types* and a *merge operator*. The calculus is inspired by two closely related calculi by Dunfield (2014) and Oliveira et al. (2016). Although Dunfield proposes a direct small-step semantics for her calculus, her semantics lacks both *determinism* and *subject-reduction*. Using our TDOS we obtain a direct semantics for $\lambda_i^{:}$ that has both properties. To fully obtain determinism, the $\lambda_i^{:}$ calculus employs a disjointness restriction proposed in Oliveira et al.'s λ_i calculus. As an added benefit the TDOS approach deals with recursion in a straightforward way, unlike λ_i and subsequent calculi where recursion is problematic. To further relate $\lambda_i^{:}$ to the calculi by Dunfield and Oliveira et al. we show two results. Firstly, the semantics of $\lambda_i^{:}$ is sound with respect to Dunfield's small-step semantics. Secondly, we show that the type system of $\lambda_i^{:}$ is complete with respect to the λ_i type system. All results have been fully formalized in the Coq theorem prover.

1 Introduction

The *merge operator* was firstly introduced by Reynolds in the Forsythe language [43] over 30 years ago. It has since been studied, refined and used in some language designs by multiple researchers [2, 5, 9, 20, 23, 39]. At its essence the merge operator allows creating values that can have *multiple types* (encoded as *intersection types* [18, 41]). For example, with the merge operator, the following program is valid:

$$let\ x : \mathsf{Int}\ \&\ \mathsf{Bool}\ =\ 1\,,,\ \mathsf{True}\ in\ (x + 1,\ \mathsf{not}\ x)$$

Here the variable x has two types, expressed by the intersection type $\mathsf{Int}\,\&\,\mathsf{Bool}$. The corresponding value for x is built using the merge operator $(, ,)$. Later uses of x, such as the expression $(x + 1,\ \mathsf{not}\ x)$ can use x both as an integer or as a boolean. For this particular example, the result of executing the expression is the pair $(2,\ \mathsf{False})$.

The merge operator adds expressive power to calculi with intersection types. Much work on intersection types has focused on *refinement intersections* [21, 24, 28], which only increase the expressive power of types. In systems with refinement intersections, types can simply be erased during compilation. However, in those systems the intersection type $\mathsf{Int}\,\&\,\mathsf{Bool}$ is invalid since Int and Bool are not refinements of each other. In other systems, including many OO languages with intersection types [25, 34, 36, 42], the type $\mathsf{Int}\,\&\,\mathsf{Bool}$ has no inhabitants and the simple program above is inexpressible. The merge operator adds expressiveness to terms and allows constructing values that inhabit the intersection type $\mathsf{Int}\,\&\,\mathsf{Bool}$.

There are various practical applications for the merge operator. One benefit, as Dunfield [23] argues, is that the merge operator and intersection types provide "*general mechanisms that subsume many different features*". This is important because often a new type system feature involves extending the metatheory and implementation, which can be non-trivial. If instead we provide general mechanisms that can encode such features, then adding new features will become a lot easier. Dunfield has illustrated this point by showing that *multi-field records*, *overloading* and forms of *dynamic typing* can all be easily encoded in the presence of the merge operator. More recently, the merge operator has been used in calculi with disjoint intersection types [2, 5, 20] to encode several non-trivial object-oriented features, which enable highly dynamic forms of object composition not available in current mainstream languages such as Scala or Java. These include *first-class traits* [4], *dynamic mixins* [2], and forms of *family polymorphism* [5]. These features allow, for instance, capturing widely used and expressive techniques for object composition used by JavaScript programmers (and programmers in other dynamically typed languages), but in a completely *statically type-safe* manner [2,4]. For example, in the SEDEL language [4], which is based on disjoint intersection types, we can define and use first-class traits such as:

```
// addId takes a trait as an argument, and returns another trait
addId(super : Trait[Person], idNumber : Int) : Trait[Student] =
  trait inherits super ⇒ {  // dynamically inheriting from an unknown person
      def id : Int = idNumber
  }
```

Similarly to classes in JavaScript, first-class traits can be passed as arguments, returned as results, and they can be constructed dynamically (at run-time). In the program above inheritance is encoded as a merge in the core language with disjoint intersection types used by SEDEL.

Despite over 30 years of research, the semantics of the merge operator has proved to be quite elusive. Because of its foundational importance, we would expect a simple and clear *direct* semantics to exist for calculi with a merge operator. After all, this is what we get for other foundational calculi such as the *simply typed lambda calculus*, *System F*, *System F_ω*, the *calculus of constructions*, *System $F_{<:}$*, *Featherweight Java* and others. All these calculi have a simple and elegant *direct operational semantics* (often presented in a small-step style [55]). While for the merge operator there have been efforts in the past to define a direct operational semantics, these efforts have placed severe limitations that disallow many of the applications previously discussed or they lacked important properties. Reynolds [44] was the first to look at this problem, but in his calculus the merge operator is severely limited (for instance a merge of two functions is not possible). Castagna [9] studied another calculus, where only merges of functions are possible. Pierce [39] was the first to briefly consider a calculus with

an unrestricted merge operator (called *glue* in his own work). He discussed an extension to $F_\wedge$ with a merge operator but he did not study the dynamic semantics with the extension. Finally, Dunfield [23] goes further and presents a direct operational semantics for a calculus with an unrestricted merge operator. However the problem is that subject-reduction and determinism are lost.

Dunfield also presents an alternative way to give the semantics for a calculus with the merge operator *indirectly by elaboration* to another calculus. This elaboration semantics is type-safe and offers, for instance, a reasonable implementation strategy, and it is also employed in more recent work on the merge operator with disjoint intersection types. However the elaboration semantics has two important drawbacks. Firstly, reasoning about the elaboration semantics is much more complex: to understand the semantics of programs with the merge operator we have to understand the translation and semantics of the target calculus. This complicates informal and formal reasoning. Secondly, a fundamental property in an elaboration semantics is *coherence* [44] (which ensures that the meaning of a program is not ambiguous). All existing calculi with disjoint intersection types prove coherence, but this currently comes at a high price: the calculi and proof techniques employed to prove coherence can only deal with terminating programs. A severe limitation in practice!

This paper presents a *type-directed operational semantics* (TDOS) for $\lambda_i^{:}$: a calculus with intersection types and a merge operator [43]. $\lambda_i^{:}$ is inspired by closely related calculi by Dunfield [23] and Oliveira et al. (λ_i) [20], but addresses two key difficulties in the dynamic semantics of calculi with a merge operator. The first one is the type-dependent nature of the merge operator. This difficulty is addressed by using types in the TDOS to guide reduction, which is crucial to prove *subject-reduction*. The second difficulty is that a fully unrestricted merge operator is inherently ambiguous. For instance the merge $1, , 2$ can evaluate to both 1 and 2. Therefore some restriction is still necessary for a deterministic semantics. To fully obtain determinism, the $\lambda_i^{:}$ calculus uses the disjointness restriction that is employed in λ_i and several other calculi using disjoint intersection types, and two important new notions: *typed reduction* and *consistency*. Typed reduction is a reduction relation that can further reduce values under a certain type. In other words, type annotations influence operational behavior: two programs that differ only in type annotations may behave differently. Consistency is an equivalence relation on values, that is key for the determinism result. Determinism in TDOS offers the same guarantee that coherence offers in an elaboration semantics (both properties ensure that the semantics is unambiguous), but it is much simpler to prove. Additionally, the TDOS approach deals with recursion in a straightforward way, unlike λ_i and subsequent calculi where recursion is very problematic for proving coherence.

To further relate $\lambda_i^{:}$ to the calculi by Dunfield and Oliveira et al. we show two results. Firstly, we show that the type system of $\lambda_i^{:}$ is complete with respect to the type system of λ_i. Secondly, the semantics of $\lambda_i^{:}$ is sound with respect to Dunfield's semantics. In our work we use two variants of $\lambda_i^{:}$: one that follows Dunfield's original formulation of subtyping, and another with a more powerful subtyping relation inspired by Bi et al. [5]. The more powerful subtyping relation enables $\lambda_i^{:}$ to account for merges of functions in a natural way, which was awkward in λ_i. For the variant with the extension we also require a minor extension to Dunfield's operational semantics. The two variants of the $\lambda_i^{:}$ calculus and its metatheory have been fully formalized in the Coq theorem prover.

In summary, the contributions of this paper are:

- **The $\lambda_i^{:}$ calculus and its TDOS:** We present a TDOS for $\lambda_i^{:}$: a calculus with intersection types and a merge operator. The semantics of $\lambda_i^{:}$ is both *deterministic* and it has *subject-reduction*.

- **Support for non-terminating programs:** Our new proof methods can deal with recursion, unlike the proof methods used in previous calculi with disjoint intersection types [5,6], due to limitations of the current proof approaches for coherence.
- **Typed reduction and consistency:** We propose the novel notions of typed reduction and consistency, which are useful to prove determinism and subject-reduction.
- **Relation with other calculus with intersection types:** We relate $\lambda_i^{:}$ with the calculi proposed by Dunfield and Oliveira et al (λ_i). In short all programs that are accepted in λ_i can type-check with our type system, and the semantics of $\lambda_i^{:}$ is sound with respect to Dunfield's semantics.
- **Coq formalization:** All the results presented in this paper have been formalized in the Coq theorem prover and they are available in the supplementary material.

2 Overview

This section gives an overview of the type-directed operational semantics for $\lambda_i^{:}$. We first provide some background about the applications of the merge operator. Then we introduce Dunfield's untyped semantics [23], and identify its problems: the non-determinism of the semantics and the lack of subject-reduction. Dunfield's semantics is nonetheless used to guide the design of our own TDOS. We show how the TDOS of $\lambda_i^{:}$ uses type annotations to guide reduction, thus obtaining a deterministic semantics that also has the subject-reduction property.

2.1 First-Class Traits: An Application of the Merge Operator

To give an idea of the kinds of applications for calculi with a merge operator, we briefly present one existing application: *typed first-class traits* [4]. *Traits* [46] in object-oriented programming provide a model of multiple inheritance. Both traits and *mixins* [27] encapsulate a collection of related methods to be added to a class. When composing multiple traits/mixins, conflicts are dealt differently. Mixins use the order of composition to determine which implementation to pick. Traits require programmers to explicitly resolve the conflicts instead, and reject compositions with conflicts. Merges with disjoint intersection types are closely related to traits because merges with conflicts are also rejected.

Here we borrow an example from the SEDEL language [4] to demostrate how it encodes (typed) first-class traits and dynamic inheritance via the merge operator.

```
type Editor = {on_key : String → String, do_cut : String, show_help : String};
type Version = {version : String};

trait editor [self : Editor & Version] ⇒ {
  on_key(key : String) = "Pressing " ++ key;
  do_cut = self.on_key "C-x" ++ " for cutting text";
  show_help = "Version: " ++ self.version ++ " Basic usage..."
};
```

In SEDEL traits are elaborated into a core calculus with disjoint intersection types and a merge operator. A trait can be viewed as a function taking a self argument and producing a record. In this example, the record, which contains three fields, is encoded as a merge of three single field records. Because all the fields have distinct field names, the merge is disjoint and the definition is accepted. Similarly to a JavaScript class, in the trait `editor`, the `doCut` method calls the `onKey` method via the self reference and it is dynamically dispatched. What is more, traits in SEDEL have a self type annotation which is similar to Scala [36]. In this example, the type of the self reference is the intersection of two record types `Editor` and

Version. Note that `show_help` is defined in terms of an undefined `version` method. Usually, in a statically typed language like Java, an abstract method is required, making `editor` an abstract class. Instead, SEDEL encodes abstract methods via self-types. The requirements stated by the type annotation of self must be satisfied when later composing `editor` with other traits, i.e. an implementation of the method `version` should be provided.

The interesting features in SEDEL are that traits are *first-class* and inheritance can be *dynamic*. The following example illustrates such features:

```
type Spelling = {check : String};

spell_mixin (base : Trait[Editor & Version, Editor]) =
  trait [self : Editor & Version] inherits base ⇒ {
    override on_key(key : String) = "Process " ++ key ++ " on spell editor";
    check = super.on_key "C-c" ++ " for spelling check"
}
```

The above function takes a trait as an argument, and returns a trait as a result. The argument `base` is a trait of type `Trait[Editor & Version, Editor]`, where the two types denote trait requirements and functionality respectively. The trait `editor` has type `Trait[Editor & Version, Editor]`, since it requires `Editor & Version` and only provides those method specified by `Editor`. Therefore, `editor` can be used as an argument for `spell_mixin`. Note that unlike mainstream OOP languages like Java, the inherited trait (which would correspond to a superclass in Java) is *parametrized*, thus enabling *dynamic inheritance*. In SEDEL the choice of the inherited trait (i.e. the superclass) can happen at run-time, unlike in languages with static inheritance (such as Java or Scala). Finally, also note the use of the keyword `override` to override `on_key`. Without such keyword the definition of `spell_mixin` would be rejected due to a conflict (or a violation of disjointness), since `base` already provides an implementation of `on_key`. For a more detailed description of SEDEL and first-class traits we refer the reader to the work by Bi et al. [4].

2.2 Background: Dunfield's Non-Deterministic Semantics

Dunfield studied the semantics of a calculus with intersection types and a merge operator. The interesting aspect of her calculus is the merge operator, which takes two terms e_1 and e_2, of some types A and B, to create a new term that can behave both as a term of type A and as a term of type B. Intersection types and the merge operator in Dunfield's calculus are similar to pair types and pairs. Indeed, a program written with pairs that behaves identically to the program shown in Section 1 is:

$$let\ x : (\mathsf{Int}, \mathsf{Bool})\ =\ (1, \mathsf{True})\ in\ (\ \mathsf{fst}\ x + 1,\ \mathsf{not}\ (\mathsf{snd}\ x)\ )$$

However while for pairs both the introductions and eliminations are explicit, with the merge operator the eliminations (i.e. projections) are *implicit* and driven by the types of the terms. Dunfield exploits this similarity to give a type-directed elaboration semantics to her calculus. The elaboration transforms merges into pairs, intersection types into pair types and inserts the missing projections.

Syntax. The top of Figure 1 shows the syntax of Dunfield's calculus. Types include a top type Top, function types $(A \rightarrow B)$ and intersection types (written as $A \,\&\, B$). Most expressions are standard, except the merges $(E_1 \,,, E_2)$. The calculus also includes a canonical top value $\top$, and allows variables to be values. Note that the original Dunfield's calculus uses a different notation for intersection types $(A \wedge B)$, and supports union types $(A \vee B)$.

$$
\begin{array}{lll}
Type & A, B ::= \mathsf{Top} \mid A \to B \mid A \mathbin{\&} B \\
Expr & E ::= x \mid \top \mid \lambda x.E \mid E_1 \; E_2 \mid \mathbf{fix}\, x.E \mid E_1 \,,, E_2 \\
Value & V ::= x \mid \top \mid \lambda x.E \mid V_1 \,,, V_2
\end{array}
$$

$\boxed{E \rightsquigarrow E'}$ $\hspace{4cm}$ *(Operational semantics of Dunfield's calculus)*

DStep-appl
$$\frac{E_1 \rightsquigarrow E_1'}{E_1\, E_2 \rightsquigarrow E_1'\, E_2}$$

DStep-appr
$$\frac{E_2 \rightsquigarrow E_2'}{V_1\, E_2 \rightsquigarrow V_1\, E_2'}$$

DStep-beta
$$\frac{}{(\lambda x.\, E)\, V \rightsquigarrow E[x \mapsto V]}$$

DStep-fix
$$\frac{}{\mathbf{fix}\, x.\, E \rightsquigarrow E[x \mapsto \mathbf{fix}\, x.\, E]}$$

DStep-mergel
$$\frac{E_1 \rightsquigarrow E_1'}{E_1 \,,, E_2 \rightsquigarrow E_1' \,,, E_2}$$

DStep-merger
$$\frac{E_2 \rightsquigarrow E_2'}{V_1 \,,, E_2 \rightsquigarrow V_1 \,,, E_2'}$$

DStep-unmergel
$$\frac{}{E_1 \,,, E_2 \rightsquigarrow E_1}$$

DStep-unmerger
$$\frac{}{E_1 \,,, E_2 \rightsquigarrow E_2}$$

DStep-split
$$\frac{}{E \rightsquigarrow E \,,, E}$$

■ **Figure 1** The syntax and non-deterministic small-step semantics of Dunfield's calculus.

Union types are not supported by $\lambda_i^{:}$, since it is based on the λ_i calculus [20] with disjoint intersection types, which does not have unions either. For a better comparison, we adjust the syntax and omit union types in Dunfield's system.

Operational Semantics. The bottom part of Figure 1 presents the reduction rules. The interesting construct is the merge operator, as all other rules not involving the merge operator are standard call-by-value reduction rules. The reduction of a merge construct in Dunfield's calculus is quite flexible: a merge of two expressions (which do not even need to be two values) can step to its left subexpression (by rule DStep-unmergel) or the right one (by rule DStep-unmerger). Any expressions can split into two by rule DStep-split. Therefore, even though the reduction rules may have already reached a value form, it is still possible to step further using rule DStep-split.

Problem 1: No Subject-Reduction. A major problem of this operational semantics is that it does not preserve types. Note that reduction is oblivious of types, so a term can reduce to two other terms with potentially different (and unrelated) types. For instance:

$$1 \,,, \mathsf{True} \rightsquigarrow 1 \qquad 1 \,,, \mathsf{True} \rightsquigarrow \mathsf{True}$$

Here the merge of an integer and a boolean is reduced to either the integer (using rule DStep-unmergel) or the boolean (using rule DStep-unmerger). In Dunfield's calculus the term $1 \,,, \mathsf{True}$ can have multiple types, including Int or Bool or even $\mathsf{Int} \mathbin{\&} \mathsf{Bool}$. As a consequence, the semantics is not type-preserving in general.

What is worse, a well-typed expression can reduce to an expression that is ill-typed:

$$(1 \,,, \lambda x.\, x + 1)\, 2 \rightsquigarrow 1\, 2$$

This reduction leads to an ill-typed term (with any type) because we drop the lambda instead of the 1 in the merge.

Problem 2: Non-determinism. Even in the case of type-preserving reductions there can be another problem. Because of the pair of unmerge rules (rule DStep-unmergel and rule DStep-unmerger), the choice between a merge always has two options. This means that a reduced term can lead to two other terms of the same type, but with different meanings. For example:

$$1 \, , , \, 2 \rightsquigarrow 1 \qquad 1 \, , , \, 2 \rightsquigarrow 2$$

There is even a third option to reduce a merge with the split rule (rule DStep-split):

$$1 \, , , \, 2 \rightsquigarrow (1 \, , , \, 2) \, , , \, (1 \, , , \, 2)$$

In other words the semantics is non-deterministic.

Note that Dunfield's operational semantics is an overapproximation of the intended behavior. In her work, it is used to provide a soundness result for her elaboration semantics, which is type-safe (but still ambiguous).

2.3 A Type-Driven Semantics for Type Preservation

An essential problem is that the semantics cannot ignore the types if the reduction is meant to be type-preserving. Dunfield notes that "*For type preservation to hold, the operational semantics would need access to the typing derivation*" [23]. To avoid run-time type-checking, we design a type-driven semantics and use type annotations to guide reduction. Therefore our $\lambda_i^{:}$ calculus is explicitly typed, unlike Dunfield's calculus. Nevertheless, it is easy to design source languages that infer some of the type annotations and insert them automatically to create valid $\lambda_i^{:}$ terms as we will see in Section 5. We discuss the main challenges and key ideas of the design of $\lambda_i^{:}$ next.

Type-driven Reduction. Our operational semantics follows a standard call-by-value small-step reduction and it is closely related to Dunfield's semantics. However, type annotations play an important role in the reduction rules and are used to guide reduction. For example, in $\lambda_i^{:}$ we can write explicitly annotated expressions such as $(1 \, , , \, \mathsf{True}) : \mathsf{Int}$ and $(1 \, , , \, \mathsf{True}) : \mathsf{Bool}$. For those expressions the following reductions are valid:

$$(1 \, , , \, \mathsf{True}) : \mathsf{Int} \hookrightarrow 1 \qquad (1 \, , , \, \mathsf{True}) : \mathsf{Bool} \hookrightarrow \mathsf{True}$$

In contrast the following reductions are not possible:

$$(1 \, , , \, \mathsf{True}) : \mathsf{Bool} \not\hookrightarrow 1 \qquad (1 \, , , \, \mathsf{True}) : \mathsf{Int} \not\hookrightarrow \mathsf{True}$$

Note also that in $\lambda_i^{:}$ the meaning of expression $1 \, , , \, \mathsf{True}$ without any type annotation can only be a corresponding value $1 \, , , \, \mathsf{True}$ that does not drop any information.

Typed Reduction. The crucial component in the operational semantics that enables the use of type information during reduction is an auxiliary *typed reduction* relation $v \hookrightarrow_A v'$ that is used when we want some value to match a type. Typed reduction is where type information from annotations in $\lambda_i^{:}$ "filters" reductions that are invalid due to a type mismatch. Typed reduction takes a value and a type (which can be viewed as inputs), and gives a unique value of that type as output. Note that this process may result in further reduction of the value, unlike many other languages where values can never be further reduced. Typed reduction is used in two places during reduction:

STEP-ANNOV
$$\frac{v \hookrightarrow_A v'}{v : A \;\hookrightarrow\; v'}$$

STEP-BETA
$$\frac{v \hookrightarrow_A v'}{(\lambda x.\, e : A \to B)\, v \;\hookrightarrow\; (e[x \mapsto v']) : B}$$

The first place where typed reduction is used is in rule STEP-ANNOV. When reduction encounters a value with a type annotation A it uses typed reduction to do further reduction depending on the type A. To see typed reduction in action, consider a simple merge of primitive values such as $1\,,,\,\mathsf{True}\,,,\,\text{`}c\text{'}$ with an annotation $\mathsf{Int}\,\&\,\mathsf{Char}$. Using rule STEP-ANNOV typed reduction is invoked, resulting in:

$$1\,,,\,\mathsf{True}\,,,\,\text{`}c\text{'} \hookrightarrow_{\mathsf{Int}\,\&\,\mathsf{Char}} 1\,,,\,\text{`}c\text{'}$$

We could have type-reduced the same value under a similar type but where the two types in the intersection are interchanged:

$$1\,,,\,\mathsf{True}\,,,\,\text{`}c\text{'} \hookrightarrow_{\mathsf{Char}\,\&\,\mathsf{Int}} \text{`}c\text{'}\,,,\,1$$

Both typed reductions are valid and they illustrate the ability of typed reduction to create a value that matches exactly with the shape of the type.

The second place where typed reduction is used is in rule STEP-BETA. In a function application, the actual argument could be a merge containing more components than what the function expects. One example is $(\lambda x.\, x + 1 : \mathsf{Int} \to \mathsf{Int})\,(1\,,,\,\mathsf{True})$. Since the merge term $(1\,,,\,\mathsf{True})$ provides an integer 1, the redundant components (the True in this case) are useless, and sometimes even harmful. Consider a function $\lambda x.\,(x\,,,\,\mathsf{False})$ with type $\mathsf{Int} \to \mathsf{Int}\,\&\,\mathsf{Bool}$, applied to $(1\,,,\,\mathsf{True})$. If we performed direct substitution of the argument in the lambda body, this would result in $1\,,,\,\mathsf{True}\,,,\,\mathsf{False}$. This brings ambiguity, and the term is not well-typed, as we shall see in Section 2.5. Therefore, before substitution, the value must be further reduced with typed reduction under the expected type of the function argument. Thus the value that is substituted in the lambda body is 1 (but not $1\,,,\,\mathsf{True}$), and the final result is $1\,,,\,\mathsf{False}$.

These examples show some non-trivial aspects of typed reduction, which must decompose values, and possibly drop some of the components and permute other components. The details of the typed reduction relation will be discussed in Section 4. As we shall see functions introduce further complications.

2.4   The Challenges of Functions

One of the hardest challenges in designing the semantics of $\lambda_i^{:}$ was the design of the rules for functions. We discuss the challenges next.

Return Types Matter.   As we have seen above, the input type annotation of lambdas is necessary during beta reduction. However, it is not enough to distinguish among multiple functions in a merge (e.g. $(\lambda x.\, x + 1)\,,,\,(\lambda x.\,\mathsf{True}))$ without run-time type checking. Unlike primitive values, whose types can be told by their forms, for functions, we need the type of the function (including the output type) to select the right function from a merge. Therefore, in $\lambda_i^{:}$ all functions are annotated with both the input and output types. With such annotations we can deal with programs such as:

$$((\lambda f.f\; 1) : (\mathsf{Int} \to \mathsf{Int}) \to \mathsf{Int})\,((\lambda x.\, x + 1) : \mathsf{Int} \to \mathsf{Int}\,,,\,(\lambda x.\,\mathsf{True}) : \mathsf{Int} \to \mathsf{Bool})$$

In this program we have a lambda that takes a function f as an argument and applies it to
1. The lambda is applied to the merge of two functions of types $\mathsf{Int} \to \mathsf{Int}$ and $\mathsf{Int} \to \mathsf{Bool}$.
To select the right function from the merge, the types of the functions are used to guide the
reduction of the merge. This avoids the need for run-time type-checking, which would be
otherwise necessary to recover the full type of functions.

Annotation Refinement.   Given a value, for any of its supertypes, typed reduction gives
a result. Since functions are values, sometimes this leads to the refinement of the type
annotation of lambdas. Following the convention introduced by previous works [20], $\to$ has
lower precedence than &, which means $A \to B \,\&\, C$ equals to $A \to (B \,\&\, C)$. Consider a single
function $\lambda x.\, x\,,\,\mathsf{True} : \mathsf{Int} \to \mathsf{Int} \,\&\, \mathsf{Bool}$ to be reduced under type $\mathsf{Int} \,\&\, \mathsf{Bool} \to \mathsf{Int}$. To let
the function return an integer when applied to a merge of type $\mathsf{Int} \,\&\, \mathsf{Bool}$, we must change
either the lambda body or the embedded annotation. Since reducing under a lambda body
is not allowed in call-by-value, $\lambda_i^{:}$ adopts the latter option, and treats the input and output
annotations differently. The input annotation should not be changed as it represents the
expected input type of the function and helps to adjust the input value before substitution
in beta reduction. The output annotation, in contrast, must be replaced by Int, representing
a future reduction to be done after substitution. The output of the application then can be
thought as an integer and can be safely merged with another boolean, for example. In short,
the actual $\lambda_i^{:}$ reduction is:

$$((\lambda x.\, x\,,\,\mathsf{True}) : \mathsf{Int} \to \mathsf{Int} \,\&\, \mathsf{Bool}) : \mathsf{Int} \,\&\, \mathsf{Bool} \to \mathsf{Int}$$
$$\hookrightarrow (\lambda x.\, x\,,\,\mathsf{True}) : \mathsf{Int} \to \mathsf{Int}$$

2.5   Disjoint Intersection Types and Consistency for Determinism

Even if the semantics is type-directed and it rules out reductions that do not preserve types,
it can still be non-deterministic. To solve this problem, we employ the disjointness restriction
proposed by Oliveira et al. [20] and the novel notion of *consistency*. Both disjointness and
consistency play a fundamental role in the proof of determinism.

Disjointness.   Two types are disjoint (written as $A * B$), if any common supertypes that
they have are *top-like types* (i.e. supertypes of any type; written as $\rceil C \lceil$).

▶ **Definition 1** (Disjoint Types)**.**

$$A * B \equiv \forall C, if\ A <: C\ and\ B <: C\ then\ \rceil C \lceil$$

Intuitively, if two types are disjoint (e.g. $\mathsf{Int} \,\&\, \mathsf{Char} * \mathsf{Bool}$), their corresponding values do not
overlap (e.g. $1\,,\,$ 'c' and True). The only exceptions are top-like types, as they are disjoint
with any types [2]. Since every value of a top-like type has the same effect, typed reduction
unifies them to a fixed result. Thus the disjointness checking in the following typing rule
guarantees that e_1 and e_2 can be merged safely, without any ambiguities. For example, this
typing rule does not accept $1\,,\,2$ or $\mathsf{True}\,,\,1\,,\,\mathsf{False}$, as two subterms of the merge have
overlapped types (in this case, the same type Int and Bool, respectively).

$$\frac{\Gamma \vdash e_1 : A \qquad \Gamma \vdash e_2 : B \qquad A * B}{\Gamma \vdash e_1\,,\,e_2 : A \,\&\, B}\ \text{\textsc{Etyp-merge}}$$

Consistency.   Recall the split rule (rule DSTEP-SPLIT) in Dunfield's semantics: $E \rightsquigarrow E ,, E$. It duplicates terms in a merge. Similar things can happen in our typed reduction if the type has overlapping parts, which is allowed, for example, in an expression $1 : \mathsf{Int} \,\&\, \mathsf{Int}$. Note that in this expression the term 1 can be given type annotation $\mathsf{Int} \,\&\, \mathsf{Int}$ since $\mathsf{Int} <: \mathsf{Int} \,\&\, \mathsf{Int}$. During reduction, typed reduction is eventually used to create a value that matches the shape of type $\mathsf{Int} \,\&\, \mathsf{Int}$ by duplicating the integer:

$$1 \hookrightarrow_{Int\&Int} 1 ,, 1$$

Note that the disjointness restriction does not allow sub-expressions in a merge to have the same type: $1 ,, 1$ cannot type-check with rule ETYP-MERGE. To obtain *type preservation*, there is a special (run-time) typing rule for merges of values, where a novel consistency check is used (written as $v_1 \approx v_2$):

$$\frac{\cdot \vdash v_1 : A \qquad \cdot \vdash v_2 : B \qquad v_1 \approx v_2}{\Gamma \vdash v_1 ,, v_2 : A \,\&\, B} \text{ ETYP-MERGEV}$$

Mainly, consistency allows values to have overlapped parts as far as they are syntactically equal. For example, $1 ,, \mathsf{True}$ and $1 ,, \text{`}c\text{'}$ are consistent, since the overlapped part $\mathsf{Int}$ in both of merges is the same value. $\mathsf{True}$ and `c` are consistent because they are not overlapped at all. But $1 ,, \mathsf{True}$ and 2 are *not consistent*, as they have different values for the same type Int. When two values have disjoint types, they must be consistent. For merges of such values, both rule ETYP-MERGEV and rule ETYP-MERGE can be applied, and the types always coincide. In $\lambda_i^{:}$, consistency is defined in terms of typed reduction:

▶ **Definition 2** (Consistency). *Two values v_1 and v_2 are said to be consistent (written $v_1 \approx v_2$) if, for any type A, the result of typed reduction for the two values is the same.*

$$v_1 \approx v_2 \equiv \forall\, A, \; if \; v_1 \hookrightarrow_A v_1' \; and \; v_2 \hookrightarrow_A v_2' \; then \; v_1' = v_2'$$

Although the specification of consistency is decidable and an equivalent algorithmic definition exists, it is not required. In practice, in a programming language implementation, the rule ETYP-MERGEV may be omitted, since, as stated, its main purpose is to ensure that run-time values are type-preserving.

Finally, note that λ_i [20] is stricter than $\lambda_i^{:}$ and forbids any intersection types which are not disjoint. That is to say, the term $1 : \mathsf{Int} \,\&\, \mathsf{Int}$ is not well-typed because the intersection $\mathsf{Int} \,\&\, \mathsf{Int}$ is not disjoint. The idea of allowing unrestricted intersections, while only having the disjointness restriction for merges, was first employed in the NeColus calculus [5]. $\lambda_i^{:}$ follows such an idea and $1 : \mathsf{Int} \,\&\, \mathsf{Int}$ is well-typed in $\lambda_i^{:}$. Allowing unrestricted intersections adds extra expressive power. For instance, in calculi with polymorphism, unrestricted intersections can be used to encode *bounded quantification* [8], whereas with disjoint intersections only such an encoding does not work [6].

3   The $\lambda_i^{:}$ Calculus: Syntax, Subtyping and Typing

This section presents the syntax, subtyping, and typing of $\lambda_i^{:}$: a calculus with intersection types and a merge operator. This calculus is a small variant of the λ_i calculus [20] (which itself is inspired by Dunfield's calculus [23]) extended with *annotated expressions*, *explicit subsumption* and *fixpoints*. The explicit type annotations and subtyping are necessary for the type-directed operational semantics of $\lambda_i^{:}$ and to preserve determinism. The addition of fixpoints illustrates the ability of TDOS to deal with non-terminating programs, which are still not supported by calculi that rely on elaboration and semantic coherence proofs [5,6].

3.1 Syntax

The syntax of $\lambda_i^{:}$ is:

$$
\begin{array}{lll}
Type & A, B ::= \mathsf{Int} \mid \mathsf{Top} \mid A \rightarrow B \mid A \,\&\, B \\
Expr & e ::= x \mid i \mid \top \mid e : A \mid e_1 \, e_2 \mid \lambda x.e : A \rightarrow B \mid e_1 \,,, e_2 \mid \mathbf{fix}\; x.e : A \\
Value & v ::= i \mid \top \mid \lambda x.e : A \rightarrow B \mid v_1 \,,, v_2 \\
Context & \Gamma ::= \cdot \mid \Gamma, x : A
\end{array}
$$

Types. Meta-variables A and B range over types. Two basic types are included: the integer type Int and the top type Top. Function types $A \rightarrow B$ and intersection types $A \,\&\, B$ can be used to construct compound types.

Expressions. Meta-variable e ranges over expressions. Expressions include some standard constructs: variables (x); integers (i); a canonical top value $\top$; annotated expressions $(e : A)$; and application of a term e_1 to term e_2 (denoted by $e_1 e_2$). Lambda abstractions $(\lambda x.e : A \rightarrow B)$ must have a type annotation $A \rightarrow B$, meaning that the input type is A and the output type is B. The expression $e_1 \,,, e_2$ is the merge of expressions e_1 and e_2. Finally, fixpoints $\mathbf{fix}\; x.\, e : A$ (which also require a type annotation) model recursion.

Values and Contexts. The meta-variable v ranges over values. Values include integers, the canonical $\top$ value, lambda abstractions and merges of values. Typing contexts are standard. Γ tracks the bound variables x with their type A.

3.2 Subtyping and Disjointness

The subtyping rules of the form $A <: B$ are shown on the top of Figure 2. These subtyping rules, except for rule S-TOPARR, were first introduced by Davies and Pfenning [21], and are used in λ_i as well. The original subtyping relation is known to be reflexive and transitive [21]. We proved the reflexivity and transitivity of the extended subtyping relation as well. There are 3 rules regarding intersection types. Together they define $A \,\&\, B$ as the greatest lower bound of A and B.

Top-like Types and Arrow Types. Intuitively, a top-like type is both a supertype and a subtype of Top, including the Top type and intersections of top-like types. The newly added rule S-TOPARR enlarges top-like types to include arrow types when their return types are top-like. A simple unary relation that captures top-like types inductively is defined on the bottom of Figure 2. The following theorem states the correctness and completeness of the definition.

▶ **Lemma 3** (Soundness and Completeness of the Definition of Top-like Types)**.**

$\rceil A \lceil$ *if and only if* $\mathsf{Top} <: A$

Rule S-TOPARR is inspired by the following rule in BCD-style subtyping [3] (and adopted by Bi et al. [5]):

$$
\frac{}{\mathsf{Top} <: \mathsf{Top} \rightarrow \mathsf{Top}} \; \text{BCD-TOPARR}
$$

$\boxed{A <: B}$ *(Subtyping)*

S-z

$$\overline{\mathsf{Int} <: \mathsf{Int}}$$

S-top

$$\overline{A <: \mathsf{Top}}$$

S-arr

$$\frac{B_1 <: A_1 \qquad A_2 <: B_2}{A_1 \to A_2 <: B_1 \to B_2}$$

S-andr

$$\frac{A_1 <: A_2 \qquad A_1 <: A_3}{A_1 <: A_2 \,\&\, A_3}$$

S-toparr

$$\frac{\mathsf{Top} <: B_2}{A <: B_1 \to B_2}$$

S-andl1

$$\frac{A_1 <: A_3}{A_1 \,\&\, A_2 <: A_3}$$

S-andl2

$$\frac{A_2 <: A_3}{A_1 \,\&\, A_2 <: A_3}$$

$\boxed{\rceil A \lceil}$ *(Top-like types)*

TL-and

$$\frac{\rceil A \lceil \qquad \rceil B \lceil}{\rceil A \,\&\, B \lceil}$$

TL-top

$$\overline{\rceil \mathsf{Top} \lceil}$$

TL-arr

$$\frac{\rceil B \lceil}{\rceil A \to B \lceil}$$

■ **Figure 2** Subtyping rules of $\lambda_i^{:}$ and definition of top-like types.

Since BCD-style subtyping includes a transitivity rule as an axiom, with this rule, $\mathsf{Int} \to \mathsf{Top}$ and $\mathsf{Int} \to (\mathsf{Top} \to \mathsf{Top})$ are supertypes (and also subtypes) of Top. Due to the lack of built-in transitivity rule in $\lambda_i^{:}$'s subtyping, the above consequence has to be expressed more explicitly in the adapted rule S-toparr. We will come back to our motivation for including rule S-toparr in Section 3.3.

Disjointness.   In Section 2.5, the specification of disjointness is presented. Such specification is a slightly more liberal version of the definition originally used in λ_i. In particular in our definition A and B themselves can be *top-like types*, which was forbidden in λ_i. An equivalent algorithmic definition of disjointness $(A *_a B)$ is presented in Appendix A, which is the same as the definition in the NeColus calculus [5].

▶ **Lemma 4** (Disjointness Properties). *Disjointness satisfies:*
1. $A * B$ *if and only if* $A *_a B$.
2. *if* $A * (B_1 \to C)$ *then* $A * (B_2 \to C)$.
3. *if* $A * B \,\&\, C$ *then* $A * B$ *and* $A * C$.

3.3   Typing

The expression typing judgment $\Gamma \vdash e : A$ is standard. It says that in the typing environment Γ the expression e is of type A. Unlike λ_i, there is no well-formedness restriction on types[1]. This generalization is inspired by the calculus NeColus [5], where the well-formedness constraints are removed from λ_i, and expressions like $1 : \mathsf{Int} \,\&\, \mathsf{Int}$ are allowed. In other words the calculus supports *unrestricted intersections* as well as *disjoint intersection types* (which are the only kind of intersections supported in λ_i).

The type system, shown in Figure 3, is syntax-directed. Most typing rules directly follow the declarative type system of λ_i, including the merge rule Etyp-merge, where disjointness is used. When two expressions have disjoint types, any parts from each of them do not have overlapping types. Therefore, their merge does not introduce ambiguity. With this

---

[1] The wellformedness and typing rules for λ_i can be found in Section 5.2.

$$\boxed{\Gamma \vdash e : A} \qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad \textit{(Typing)}$$

ETYP-TOP

$$\overline{\Gamma \vdash \top : \mathsf{Top}}$$

ETYP-LIT

$$\overline{\Gamma \vdash i : \mathsf{Int}}$$

ETYP-VAR

$$\frac{x : A \in \Gamma}{\Gamma \vdash x : A}$$

ETYP-ANNO

$$\frac{\Gamma \vdash e : B \qquad B <: A}{\Gamma \vdash (e : A) : A}$$

ETYP-ABS

$$\frac{\Gamma, x : A \vdash e : B \qquad C <: A \qquad B <: D}{\Gamma \vdash (\lambda x.\, e : A \to D) : C \to D}$$

ETYP-FIX

$$\frac{\Gamma, x : A \vdash e : A}{\Gamma \vdash (\mathbf{fix}\, x.\, e : A) : A}$$

ETYP-APP

$$\frac{\Gamma \vdash e_1 : A \to B \qquad \Gamma \vdash e_2 : A}{\Gamma \vdash e_1\, e_2 : B}$$

ETYP-MERGE

$$\frac{\Gamma \vdash e_1 : A \qquad \Gamma \vdash e_2 : B \qquad A * B}{\Gamma \vdash e_1\, , , e_2 : A \,\&\, B}$$

ETYP-MERGEV

$$\frac{\cdot \vdash v_1 : A \qquad \cdot \vdash v_2 : B \qquad v_1 \approx v_2}{\Gamma \vdash v_1\, , , v_2 : A \,\&\, B}$$

■ **Figure 3** Type system of $\lambda_i^{:}$.

restriction, rule ETYP-MERGE does not accept expressions like $1\, , , 2$ or even $1\, , , 1$. On the other hand, the novel rule ETYP-MERGEV allows *consistent* values to be merged regardless of their types. It accepts $1\, , , 1$ while still rejecting $1\, , , 2$. It is for values only, and values are closed. Therefore the type judgments appearing in it as premises should have empty context, which is denoted by $\cdot$. Together the two rules support the determinism and type preservation of the TDOS, as discussed in Section 2.5.

Top-Like Types and Merges of Functions. We can finally come back to the motivation to include rule S-TOPARR in subtyping and depart from both Dunfield calculus and λ_i, which do not have such a rule. *Without the rule* S-TOPARR *in subtyping*, no arrow types are top-like, therefore two arrow types $A \to B$ and $C \to D$ are never disjoint in terms of Definition 1, as they have a common supertype $A \,\&\, C \to \mathsf{Top}$. Consequently, we can never create merges with more than one function, which is quite restrictive. For Dunfield this is not a problem, because she does not have the disjointness restriction. So her calculus supports merges of any functions (but it is incoherent). In λ_i an ad-hoc solution is proposed, by forcing the matter and employing the syntactic definition of top-like types in Figure 2 in disjointness. However this means that in λ_i Lemma 3 is false, since top-like function types are not supertypes of Top. In contrast, the approach we take in $\lambda_i^{:}$ is to add the rule S-TOPARR in subtyping. Now $\mathsf{Top} <: (A \,\&\, C \to \mathsf{Top})$ is derivable and thus $A \,\&\, C \to \mathsf{Top}$ is genuinely a top-like type. In turn this makes merges of multiple functions typeable without losing the intuition behind top-like types.

Type-Checking for Lambda Abstractions. Rule ETYP-ABS can be thought as a combination of the standard typing rule and the subsumption rule. A well-typed lambda abstraction can have multiple types with the same return type. Its type annotation indicates the *principal input type* and the return type. Thus the input type can be any subtype of the principal one, since arrow types are contravariant in their argument types. While the principal input type describes the lambda's expectation on its argument, the annotated return type ensures the type of the evaluated result of lambdas. It just needs to be a supertype of the inner expression of the lambda. Rule ETYP-ABS is inspired by the "distributed" use of subsumption in the $\lambda\&$ calculus [9].

Explicit Subsumption.   Unlike many calculi where there is a general subsumption rule that can apply anywhere, in $\lambda_i^{:}$ subsumption needs to be explicitly triggered by a type annotation (except for lambdas, as explained above). The annotation rule ETYP-ANNO acts as explicit subsumption and assigns supertypes to expressions, provided a suitable type annotation. There is a strong motivation not to include a general (implicit) subsumption rule in calculi with disjoint intersection types. With an implicit subsumption rule disjointness alone is insufficient to prevent some ambiguous terms, as shown in the following example.

$$
\text{Subsumption} \; \dfrac{
\text{Etyp-merge} \; \dfrac{
\text{Subsumption} \; \dfrac{\cdot \vdash 1 : \mathsf{Int} \qquad \mathsf{Int} <: \mathsf{Top}}{\cdot \vdash 1 : \mathsf{Top}} \qquad \cdot \vdash 2 : \mathsf{Int} \qquad \mathsf{Top} * \mathsf{Int}
}{\cdot \vdash 1\,,,2 : \mathsf{Top}\,\&\,\mathsf{Int}} \qquad \mathsf{Top}\,\&\,\mathsf{Int} <: \mathsf{Int}\,\&\,\mathsf{Int}
}{\cdot \vdash 1\,,,2 : \mathsf{Int}\,\&\,\mathsf{Int}}
$$

Via the typical implicit subsumption, type Top is assigned to integer 1. Then 1 can be merged with 2 of type Int since their types are disjoint. At that time, the merged term $1\,,,2$ has type $\mathsf{Top}\,\&\,\mathsf{Int}$, which is a subtype of $\mathsf{Int}\,\&\,\mathsf{Int}$. By applying the subsumption rule again, the ambiguous term $1\,,,2$ finally bypasses the disjointness restriction, having type $\mathsf{Int}\,\&\,\mathsf{Int}$. However, note that with rule ETYP-ANNO we can still type-check the term $(1 : \mathsf{Top})\,,,2$, and reducing that term under the type *Int* can only unambiguously result in 2. The type annotation is key to prevent using the value 1 as an integer. Finally, the use of an explicit subsumption rule is a simpler alternative to bidirectional type-systems employed in other calculi with disjoint intersection types. Bidirectional type-checking is also capable of controlling subsumption, but adds more complexity.

Principal Types.   The principal type of a value is the most specific one among all of its types, i.e. it is the subtype of any other type of the term. Its definition is syntax-directed.

▶ **Definition 5** (Principal types). *$type_p\langle v \rangle$ calculates the principal type of value v.*

$$
\begin{aligned}
type_p\langle \top \rangle &= \mathsf{Top} \\
type_p\langle n \rangle &= \mathsf{Int} \\
type_p\langle \lambda x.\, e : A \to B \rangle &= A \to B \\
type_p\langle v_1\,,,v_2 \rangle &= type_p\langle v_1 \rangle \; \& \; type_p\langle v_2 \rangle
\end{aligned}
$$

▶ **Lemma 6** (Principal Types). *For any value v, if its principal type is A, then*
1. *if $\cdot \vdash v : B$ then $A <: B$.*
2. *if $\cdot \vdash v : B$ and $B * C$ then $A * C$.*
3. *$\cdot \vdash v : A$.*

4   A Type-Directed Operational Semantics for $\lambda_i^{:}$

This section introduces the type-directed operational semantics for $\lambda_i^{:}$. The operational semantics uses type information arising from type annotations to guide the reduction process. In particular, a new relation called *typed reduction* is used to further reduce values based on the contextual type information, forcing the value to match the type structure. We show two important properties for $\lambda_i^{:}$: *determinism of reduction* and *type soundness*. That is to say, there is only one way to reduce an expression according to the small-step relation, and the process preserves types and never gets stuck.
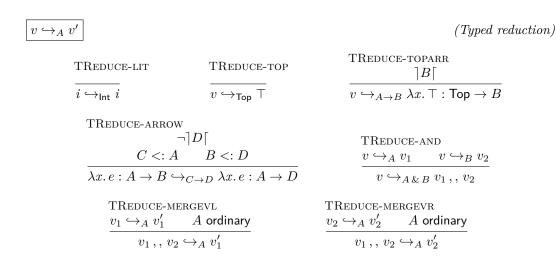
$$\boxed{v \hookrightarrow_A v'} \hspace{6cm} \textit{(Typed reduction)}$$

$$\frac{}{\text{TReduce-lit}} \qquad \frac{}{\text{TReduce-top}}$$

$$\frac{}{i \hookrightarrow_{\mathsf{Int}} i} \qquad \frac{}{v \hookrightarrow_{\mathsf{Top}} \top}$$

$$\text{TReduce-toparr}$$
$$\frac{\lceil B \rceil}{v \hookrightarrow_{A \to B} \lambda x.\top : \mathsf{Top} \to B}$$

$$\text{TReduce-arrow}$$
$$\frac{\neg \lceil D \rceil \qquad C <: A \qquad B <: D}{\lambda x.\, e : A \to B \hookrightarrow_{C \to D} \lambda x.\, e : A \to D}$$

$$\text{TReduce-and}$$
$$\frac{v \hookrightarrow_A v_1 \qquad v \hookrightarrow_B v_2}{v \hookrightarrow_{A\,\&\,B} v_1 \,,,\, v_2}$$

$$\text{TReduce-mergevl}$$
$$\frac{v_1 \hookrightarrow_A v_1' \qquad A \text{ ordinary}}{v_1 \,,,\, v_2 \hookrightarrow_A v_1'}$$

$$\text{TReduce-mergevr}$$
$$\frac{v_2 \hookrightarrow_A v_2' \qquad A \text{ ordinary}}{v_1 \,,,\, v_2 \hookrightarrow_A v_2'}$$

■ **Figure 4** Typed reduction of $\lambda_i^{\;:}$.

4.1 Typed Reduction of Values

To account for the type information during reduction $\lambda_i^{\;:}$ uses an auxiliary reduction relation called *typed reduction* for reducing values under a certain type. Typed reduction $v \hookrightarrow_A v'$ reduces the value v under type A, producing a value v' that has type A. It arises when given a value v of some type, where A is a supertype of the type of v, and v needs to be converted to a value compatible with the supertype A. The typed reduction ensures that values and types have a strong correspondence. If a value is well-typed, its principal type can be told directly by looking at its syntactic form.

Figure 4 shows the typed reduction relation. Rule TReduce-top expresses the fact that Top is the supertype of any type, which means that any value can be reduced under type Top. Similarly, rule TReduce-toparr indicates that any value reduces to a lambda abstraction $\lambda x.\top : \mathsf{Top} \to B$ under a top-like arrow type $A \to B$. Although it is not the only inhabited value of type $A \to B$, the reduction result has to be fixed for determinism. Rule TReduce-lit expresses that an integer value reduced under the supertype Int is just the integer value itself. Rule TReduce-arrow states that a lambda value $\lambda x.\, e : A \to B$, under a *non-top-like type* $C \to D$, evaluates to $\lambda x.\, e : A \to D$ if $C <: A$ and $B <: D$. The restriction that $C \to D$ is not top-like avoids overlapping with rule TReduce-toparr. Importantly rule TReduce-arrow changes the return type of lambda abstractions. For example:

$$(\lambda x.\, x \,,,\, 2 : \mathsf{Char} \to \mathsf{Char}\,\&\,\mathsf{Int}) \hookrightarrow_{(\mathsf{Char}\,\&\,\mathsf{Int} \to \mathsf{Char})} \lambda x.\, x \,,,\, 2 : \mathsf{Char} \to \mathsf{Char}$$

Intersections and Merges. In the remaining rules, we first decompose intersections. Then we only need to consider types that are not intersections, which are called *ordinary types* [21]:

$$\boxed{A \text{ ordinary}} \hspace{6cm} \textit{(Ordinary types)}$$

$$\text{O-top} \qquad\qquad \text{O-int} \qquad\qquad \text{O-arrow}$$

$$\frac{}{\mathsf{Top} \text{ ordinary}} \qquad \frac{}{\mathsf{Int} \text{ ordinary}} \qquad \frac{}{A \to B \text{ ordinary}}$$

We take care of the value by going through every merge, until both value and types are in a basic form. Rule TReduce-mergevl and rule TReduce-mergevr are a pair of rules for reducing merges under an ordinary type. Since the type is not an intersection, the result contains no merge. Usually, we need to select between the left part and right part of a merge according to the type. The values of disjoint types do not overlap on non-top-like types. For example, $1 \,,\, (\lambda x.\, x : \mathsf{Int} \to \mathsf{Int}) \hookrightarrow_{\mathsf{Int}} 1$ selects the left part. For top-like types, no matter which rule is applied, the reduction result is determined by the type only, as rule TReduce-top and rule TReduce-toparr suggest.

Rule TReduce-and is the rule that deals with intersection types. It says that if a value v can be reduced to v_1 under type A, and can be reduced to v_2 under type B, then its reduction result under type $A \,\&\, B$ is the merge of two results $v_1 \,,\, v_2$. Note that this rule may *duplicate values*. For example $1 \hookrightarrow_{\mathsf{Int}\,\&\,\mathsf{Int}} 1 \,,\, 1$. Such duplication requires special care, since the merge violates disjointness. The specially designed typing rule (rule ETyp-mergev) uses the notion of consistency (discussed in Section 4.2) instead of disjointness to type-check a merge of two values. Note also that such duplication implies that sometimes it is possible to use either rule TReduce-mergevl or rule TReduce-mergevr to reduce a value. For example, $1 \,,\, 1 \hookrightarrow_{\mathsf{Int}} 1$. The consistency restriction (Definition 2) in rule ETyp-mergev ensures that no matter which rule is applied in such a case, the result is the same.

Example.   A larger example to demonstrate how typed reduction works is:

$$(\lambda x.\, x \,,\, \text{`}c\text{'} : \mathsf{Int} \to \mathsf{Int}\,\&\,\mathsf{Char}) \,,\, (\lambda x.\, x : \mathsf{Bool} \to \mathsf{Bool}) \,,\, 1$$
$$\hookrightarrow_{\mathsf{Int}\,\&\,(\mathsf{Int}\to\mathsf{Int})} 1 \,,\, (\lambda x.\, x \,,\, \text{`}c\text{'} : \mathsf{Int} \to \mathsf{Int})$$

The initial value is the merge of two lambda abstractions and an integer. The target type is $\mathsf{Int}\,\&\,(\mathsf{Int}\to\mathsf{Int})$. Because the target type is an intersection, typed reduction first employs rule TReduce-and to decompose the intersection into Int and $\mathsf{Int} \to \mathsf{Int}$. Under type Int the value reduces to 1, and under type $\mathsf{Int} \to \mathsf{Int}$ it will reduce to $\lambda x.\, x \,,\, \text{`}c\text{'} : \mathsf{Int} \to \mathsf{Int}$. Therefore, we obtain the merge $1 \,,\, (\lambda x.\, x \,,\, \text{`}c\text{'} : \mathsf{Int} \to \mathsf{Int})$ with type $\mathsf{Int}\,\&\,(\mathsf{Int}\to\mathsf{Int})$.

Basic Properties of Typed Reduction.   Some properties of typed reduction can be proved directly by induction on the typed reduction derivation. First, when typed reduction is under a top-like type, the result only depends on the type. Second, typed reduction produces the same result whenever it is done directly or indirectly. Third, if a well-typed value can be typed reduced by some type, its principal type must be a subtype of that type.

▶ **Lemma 7** (Typed reduction on top-like types). *If $\rceil A \lceil$, $v_1 \hookrightarrow_A v_1'$, and $v_2 \hookrightarrow_A v_2'$ then $v_1' = v_2'$.*

When typed reduction is under a top-like type, the result only depends on the type.

▶ **Lemma 8** (Transitivity of typed reduction). *If $v \hookrightarrow_A v_1$, and $v_1 \hookrightarrow_B v_2$, then $v \hookrightarrow_B v_2$.*

Typed reduction produces the same result whenever it is done directly or indirectly.

▶ **Lemma 9** (Typed reduction respects subtyping). *If $v \hookrightarrow_A v'$, then $type_p\langle v \rangle <: A$.*

This lemma relates typed reduction and subtyping. It states that if a well-typed value can be typed reduced by type A, its principal type must be a subtype of A.

4.2   Consistency and Type Soundness of Typed Reduction

Consistent values, as specified in Definition 2, introduce no ambiguity in typed reduction. Consider one type, if two consistent values both can reduce under the type, they should produce the same result. The *consistency* restriction ensures that duplicated values in a merge type-check, but it still rejects merges with different values of the same type. A value of a top-like type is consistent with any other value. It only type reduces under top-like types, which leads to a fixed result decided by the type.

Relating Disjointness and Consistency.   Assuming that two values have disjoint types, according to Lemma 6, their principal types must be disjoint as well. From Lemma 9, we can conclude that when the two values both reduce under a type, that type must be a common supertype of their principal types, which is known to be top-like (Definition 1). Furthermore, Lemma 7 implies that their reduction results are always the same under such top-like types, so they are consistent (Definition 2).

▶ **Lemma 10** (Consistency of disjoint values). *If $A * B$, $\cdot \vdash v_1 : A$, and $\cdot \vdash v_2 : B$ then $v_1 \approx v_2$.*

Determinism and Type Soundness of Typed Reduction.   The merge construct makes it hard to design a deterministic operational semantics. Disjointness and consistency restrictions prevent merges like $1, , 2$, and bring the possibility to deal with merges based on types. Typed reduction takes a well-typed value, which, if it is a merge, must be consistent (according to Lemma 10). When the two typed reduction rules for merges (rule TReduce-mergevl and rule TReduce-mergevr) overlap, no matter which one is chosen, either value reduces to the same result due to consistency (Definition 2). Indeed our typed reduction relation always produces a unique result for any legal combination of the input value and type. This serves as a foundation for the determinism of the operational semantics.

▶ **Lemma 11** (Determinism of Typed Reduction). *For every well-typed v (that is there is some type B such that $\cdot \vdash v : B$), if $v \hookrightarrow_A v_1$ and $v \hookrightarrow_A v_2$ then $v_1 = v_2$.*

Via the transitivity lemma (Lemma 8) and the above determinism lemma, we obtain the following property: any reduction results of the given value are consistent.

▶ **Lemma 12** (Consistency after Typed Reduction). *If v is well-typed , and $v \hookrightarrow_A v_1$, and $v \hookrightarrow_B v_2$ then $v_1 \approx v_2$.*

The lemma shows that the reduction result of rule TReduce-and is always made of consistent values, which is needed in type preservation via the typing rule ETyp-mergev. Then a (generalized) type preservation lemma on typed reduction can be proved.

▶ **Lemma 13** (Preservation of Typed reduction). *If $\cdot \vdash v : B$ and $v \hookrightarrow_A v'$ then $\cdot \vdash v' : A$.*

In the particular case where $A = B$, this lemma shows that typed reduction preserves types. However, more generally, it shows that if a value is well-typed under a type B and it can be type reduced under another type A then the reduced value is always well-typed at type A. Finally, the typed reduction progress lemma is:

▶ **Lemma 14** (Progress of Typed Reduction). *If $\cdot \vdash v : A$, and $A <: B$, then $\exists v'$, $v \hookrightarrow_B v'$.*

$$\boxed{e \hookrightarrow e'} \qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad \textit{(Reduction)}$$

$$
\frac{\text{STEP-APPL}}{\dfrac{e_1 \hookrightarrow e_1'}{e_1\, e_2 \hookrightarrow e_1'\, e_2}}
\qquad
\frac{\text{STEP-APPR}}{\dfrac{e_2 \hookrightarrow e_2'}{v_1\, e_2 \hookrightarrow v_1\, e_2'}}
\qquad
\frac{\text{STEP-FIX}}{\mathbf{fix}\, x.\, e : A \hookrightarrow e[x \mapsto \mathbf{fix}\, x.\, e : A]}
$$

$$
\frac{\text{STEP-MERGEL}}{\dfrac{e_1 \hookrightarrow e_1'}{e_1\, ,, e_2 \hookrightarrow e_1'\, ,, e_2}}
\qquad
\frac{\text{STEP-MERGER}}{\dfrac{e_2 \hookrightarrow e_2'}{v_1\, ,, e_2 \hookrightarrow v_1\, ,, e_2'}}
\qquad
\frac{\text{STEP-ANNO}}{\dfrac{e \hookrightarrow e'}{e : A \hookrightarrow e' : A}}
$$

$$
\frac{\text{STEP-BETA}}{\dfrac{v \hookrightarrow_A v'}{(\lambda x.\, e : A \to B)\, v \hookrightarrow (e[x \mapsto v']) : B}}
\qquad\qquad
\frac{\text{STEP-ANNOV}}{\dfrac{v \hookrightarrow_A v'}{v : A \hookrightarrow v'}}
$$

🟨 **Figure 5** Call-by-value reduction of $\lambda_i^{:}$.

4.3   Reduction

The reduction rules are presented in Figure 5. Most of them are standard. Rule STEP-BETA and rule STEP-ANNOV are the two rules relying on typed reduction judgments. Rule STEP-BETA says that a lambda value $\lambda x.\, e : A \to B$ applied to value v reduces by replacing the bound variable x in e by v'. Importantly v' is obtained by type reducing v under type A. In other words, in rule STEP-BETA further (typed) reduction may be necessary on the argument depending on its type. This is unlike many other calculi where values are in a final form and no further reduction is needed (thus the value v can be directly substituted). The rule STEP-ANNOV says that an annotated $v : A$ can be reduced to v' if v type reduces to v' under type A.

Metatheory of Reduction.   When designing the operational semantics of $\lambda_i^{:}$, we want it to have two properties: *determinism of reduction* and *type soundness*. That is to say, there is only one way to reduce an expression according to the small-step relation, and the process preserves types and never gets stuck. Similar lemmas on typed reduction were already presented, which are necessary for proving the following theorems, mainly in cases related to rule STEP-ANNOV and rule STEP-BETA.

▶ **Theorem 15** (Determinism of $\hookrightarrow$). *If $\cdot \vdash e : A$, $e \hookrightarrow e_1$, $e \hookrightarrow e_2$, then $e_1 = e_2$.*

▶ **Theorem 16** (Type Preservation of $\hookrightarrow$). *If $\cdot \vdash e : A$, and $e \hookrightarrow e'$ then $\cdot \vdash e' : A$.*

▶ **Theorem 17** (Progress of $\hookrightarrow$). *If $\cdot \vdash e : A$, then e is a value or $\exists e'$, $e \hookrightarrow e'$.*

5   Relationship to Dunfield's Calculus and λ_i

Dunfield's calculus [23] and λ_i [20] are two calculi that directly inspired $\lambda_i^{:}$. In this section, we discuss the relationship between $\lambda_i^{:}$ and them. First, we show that $\lambda_i^{:}$'s TDOS and a slightly extended version of Dunfield's non-deterministic operational semantics are related. The need for extending Dunfield's original semantics is mostly due to the addition of the rule S-TOPARR in subtyping, which Dunfield does not have. In Section 6 we also discuss a variant of $\lambda_i^{:}$ (which does not include rule S-TOPARR) and show that such variant requires

$$
\begin{aligned}
|\,i\,| &= i \\
|\,\top\,| &= \top \\
|\,\lambda x.\,e : A \to B\,| &= \lambda x.|\,e\,| \\
|\,\mathbf{fix}\,x.\,e : A\,| &= \mathbf{fix}\,x.|\,e\,| \\
|\,e : A\,| &= |\,e\,| \\
|\,e_1\,e_2\,| &= |\,e_1\,|\,|\,e_2\,| \\
|\,e_1\,,,\,e_2\,| &= |\,e_1\,|\,,,\,|\,e_2\,|
\end{aligned}
$$

Figure 6 Type erasure for $\lambda_i^{:}$ expressions.

no changes to Dunfield's original semantics. The other relationship is between $\lambda_i^{:}$'s type system and λ_i's type system. The former comparison shows the soundness of the operational semantics of $\lambda_i^{:}$ with respect to Dunfield's semantics. The latter one proves that $\lambda_i^{:}$'s type system is at least as expressive as, if not stronger than, λ_i's.

Type Erasure.   Differently from the other two systems, $\lambda_i^{:}$ uses type annotations in its syntax to obtain a direct operational semantics. $|\,e\,|$ erases annotations in term e. By erasing all annotations, terms in $\lambda_i^{:}$ can be converted to terms in Dunfield's system and λ_i. The only exception is fixpoints, which λ_i does not have. The annotation erasure function is defined in Figure 6. Note that for every value v in $\lambda_i^{:}$, $|\,v\,|$ is a value as well.

5.1   Soundness with respect to Dunfield's Operational Semantics

Dunfield's original reduction rules are presented in Fig 1. We extend her operational semantics with the following two rules. The full reduction rules can be found in the appendix.

$\boxed{E \rightsquigarrow E'}$ *(The extension of Dunfield's calculus)*

$$
\frac{}{E \rightsquigarrow \top}\;\text{DStep-top}
\qquad\qquad
\frac{}{\top \rightsquigarrow \lambda x.\,\top}\;\text{DStep-toparr}
$$

Rule DStep-top states that any value can be reduced to $\top$, corresponding to $A <: \mathsf{Top}$. Rule DStep-toparr says that the value $\top$ can be reduced to a lambda which returns $\top$, suggested by the subtyping rule S-toparr. Together with merge rules, the extended reduction can reduce any term to a value under a top-like type. Dunfield avoids having a rule DStep-top by performing a simplifying elaboration step advance:

$$
\frac{}{\Gamma \vdash V : \mathsf{Top} \hookrightarrow \top}\;\text{Dunfield-Typing-T}
$$

With such a rule, values of type Top are directly translated into $\top$, and do not need any further reduction in the target language. Accordingly, in her source language, there is no rule to convert these values to $\top$. We do not have such an elaboration step and we have already added rule DStep-toparr, so instead, we extend the original semantics with the two rules.

Soundness.   Given Dunfield's extended semantics, we can show a theorem that each step in the TDOS of λ_i corresponds to zero, one, or multiple steps in Dunfield's semantics.

▶ **Theorem 18** (Soundness of $\hookrightarrow$ with respect to Dunfield's semantics). *If* $e \hookrightarrow e'$*, then* $|\,e\,| \rightsquigarrow^* |\,e'\,|$*.*

$\boxed{\Gamma \models A}$ *(Type wellformedness)*

$$
\frac{}{\Gamma \models \mathsf{Top}} \text{ Wf-top}
\qquad
\frac{}{\Gamma \models \mathsf{Int}} \text{ Wf-int}
\qquad
\frac{\Gamma \models A \qquad \Gamma \models B}{\Gamma \models A \to B} \text{ Wf-arr}
\qquad
\frac{\Gamma \models A \qquad \Gamma \models B \qquad A *_i B}{\Gamma \models A \,\&\, B} \text{ Wf-and}
$$

$\boxed{\Gamma \models E : A}$ *(Typing)*

$$
\frac{}{\Gamma \models \top : \mathsf{Top}} \text{ ITyp-top}
\qquad
\frac{}{\Gamma \models i : \mathsf{Int}} \text{ ITyp-lit}
\qquad
\frac{x : A \in \Gamma}{\Gamma \models x : A} \text{ ITyp-var}
\qquad
\frac{\Gamma \models A \qquad \Gamma, x : A \models E : B}{\Gamma \models (\lambda x.\, E) : A \to B} \text{ ITyp-lam}
$$

$$
\frac{\Gamma \models E_1 : A \to B \qquad \Gamma \models E_2 : A}{\Gamma \models E_1\, E_2 : B} \text{ ITyp-app}
\qquad
\frac{\Gamma \models E_1 : A \qquad \Gamma \models E_2 : B \qquad A *_i B}{\Gamma \models E_1 \,,,\, E_2 : A \,\&\, B} \text{ ITyp-merge}
\qquad
\frac{\Gamma \models E : A \qquad A \lessdot B}{\Gamma \models E : B} \text{ ITyp-sub}
$$

■ **Figure 7** The declarative type system of λ_i.

A necessary lemma for this theorem is the soundness of typed reduction.

▶ **Lemma 19** (Soundness of Typed Reduction with respect to Dunfield's semantics). *If $v \hookrightarrow_A v'$, then $|v| \rightsquigarrow^* |v'|$.*

This lemma shows that although the type information guides the reduction of values, it does not add additional behavior to values. For example, a merge can step to its left part (or the right part) with rule TReduce-mergevl (or rule TReduce-mergevr), corresponding to rule DStep-unmergel (or rule DStep-unmerger). And rule TReduce-and ($v \hookrightarrow_{A\,\&\,B} v_1 \,,,\, v_2$ if $v \hookrightarrow_A v_1$ and $v \hookrightarrow_B v_2$) can be understood as a combination of spliting (rule DStep-split $V \rightsquigarrow V \,,,\, V$) and further reduction on each component separately.

In Section 6, we present another variant of $\lambda_i^:$, which has the same subtyping relation as Dunfield's system (minus union types). The same soundness theorem is proved for that variant without any modifications to Dunfield's operational semantics.

5.2 Completeness with respect to the Type System of λ_i

λ_i drops union types and introduces the disjointness restriction to Dunfield's system. When introducing λ_i, Oliveira et al. proposed an algorithmic and a declarative type system. The two type systems were shown to be equally expressive. For the declarative type system there is still the possibility of ambiguity due to the presence of an (implicit) subsumption rule (see also the discussion in Section 3.3). However, annotations in the bidirectional algorithmic type system ensure that well-typed terms in λ_i are unambiguous and subsumption is kept under control.

The type system of $\lambda_i^:$ is based on the declarative type system of λ_i, with three main changes:

1. $\lambda_i^{:}$ forces the subsumption rule to be explicitly triggered by a type annotation.
2. $\lambda_i^{:}$ supports fixpoints while λ_i does not.
3. $\lambda_i^{:}$ has an additional rule for the merge of values (rule ETYP-MERGEV), which is required to prove type preservation, since duplicated values can occur in merges after reduction.

Some details need to be explained before presenting the completeness theorem. Firstly, because they are irrelevant, rules related to products and projection operators in λ_i are dropped. Secondly, the subtyping in $\lambda_i^{:}$ is stronger due to the added rule S-TOPARR. Thirdly, top-like types are disjoint with any type in $\lambda_i^{:}$, while the disjointness in λ_i is restricted to types which are not top-like. The definition of λ_i's subtyping and disjointness can be found in the appendix.

▶ **Theorem 20** (Completeness of Typing with respect to λ_i). *If $\Gamma \models E : A$, then there exists some e such that $\Gamma \vdash e : A$ and $E = |e|$.*

The above theorem shows that the type system of $\lambda_i^{:}$ is able to type check any well-typed terms in λ_i, with proper type annotations inserted based on the typing derivation. It is built on the completeness of subtyping and disjointness of $\lambda_i^{:}$. The result means that λ_i's type system (or any type system equivalent to it) can be used as a surface language where many of the explicit annotations of $\lambda_i^{:}$ are inferred automatically. That is to say, the λ_i calculus can be translated without loss of expressivity or flexibility into $\lambda_i^{:}$.

To further show that some type inference with recursion is feasible, we extended the bi-directional type system of λ_i with recursion, and replaced the subtyping and disjointness by $\lambda_i^{:}$'s. We designed an elaboration from the extended system to $\lambda_i^{:}$ and proved the following theorem. The typing rules can be found in the appendix.

▶ **Theorem 21** (Completeness of Typing with respect to the Extended Bidirectional Type System of λ_i). *If $\Gamma \vdash E \Rightarrow A \hookrightarrow e$ or $\Gamma \vdash E \Leftarrow A \hookrightarrow e$, then $\Gamma \vdash e : A$.*

6   Discussion

This section discusses one variant of $\lambda_i^{:}$, which is also formalized in Coq. The variant follows the subtyping relation in λ_i and Dunfield's calculus strictly and does not support multiple functions in merges. Some possible extensions to our work are also discussed.

6.1   A Variant of $\lambda_i^{:}$

In Section 5.1, we validate the TDOS of $\lambda_i^{:}$ via a soundness theorem (Theorem 18) with respect to an extended operational semantics of Dunfield's calculus. In this section, we discuss a variant of $\lambda_i^{:}$ that requires no extension to Dunfield's operational semantics. Its syntax and typing rules can be found in the appendix. Instead of adding rule S-TOPARR, this variant keeps the same subtyping relation as Dunfield's and adapts the definition of top-like types and disjointness, losing the ability to have multiple functions in a merge. Consequently, it is possible to prove the following soundness theorem on this variant without any modifications on Dunfield's operational semantics[2].

▶ **Theorem 22** (Soundness of $\hookrightarrow$ in the simple variant). *If $e \hookrightarrow e'$, then $|e| \rightsquigarrow^* |e'|$.*

The above theorem states that each step taken by the TDOS corresponds to a series of reduction in the original operational semantics of Dunfield's calculus.

---

[2] For the syntax and rules of Dunfield's system, please refer to Section 2.

Besides soundness, this variant keeps the other important properties of λ_i: *determinism*, *type preservation* and *progress*. A completeness theorem with respect to the type system of λ_i is established as well.

▶ **Theorem 23** (Completeness of Typing in the simple variant). *If $\Gamma \models E : A$ in λ_i, then there exists some e such that $\Gamma \vdash e : A$ in the variant and $E = \lfloor e \rfloor$.*

Designing the Variant.   As presented in Section 5.1, there are two reduction rules in $\lambda_i^:$ that are related to the extension of Dunfield's operational semantics: rule TREDUCE-TOP ($v \hookrightarrow_{\mathsf{Top}} \top$) and rule TREDUCE-TOPARR. They reduce values under top-like types into a unified form. Without rule S-TOPARR, no arrow types are top-like, thus the latter is removed from the variant. However there is still rule TREDUCE-TOP, which is not accounted for in Dunfield's original system. While we believe that such a rule fits in spirit well with the remaining non-deterministic rules, it is interesting to see if it is possible to model a calculus without it (and without extending Dunfield semantics at all).

Reducing a value v under type Top, in fact, can be thought as seeking an inhabited value of Top which acts like v. In Dunfield's original semantics there is no way to convert a value to $\top$, which is our source of difficulties. Dunfield solves this problem by having a typing rule for values that allows any value to have type Top. This works well in her setting because she can have ambiguous terms. Unfortunately, it does not work well in our setting because, as discussed in detail in Section 3.3, allowing values to be implicitly typed as Top provides a way to bypass the disjointness restrictions. We overcome the problem instead by introducing a new value construct $v : \mathsf{Top}$ in our variant. This new form of value inhabits the Top type but, unlike $\top$ it does not forget about the original value v (which can be of any type). Thus the original value v, can be recovered by erasing the wrapped annotation.

TRED-TOP

$$\frac{}{v \hookrightarrow_{\mathsf{Top}} v : \mathsf{Top}}$$

Although the new rule then corresponds to $v \rightsquigarrow v$ after annotation erasure, it breaks determinism as a merge can reduce to either its left or right component, leading to different results, e.g. $1,,\mathsf{True} \hookrightarrow_{\mathsf{Top}} 1 : \mathsf{Top}$ and $1,,\mathsf{True} \hookrightarrow_{\mathsf{Top}} \mathsf{True} : \mathsf{Top}$. To solve this problem we directly reduce the value before splitting merges by excluding Top from ordinary types.

To use the new construct for expressions and mix it with annotated terms, values and expressions are separated into two syntactic categories in the variant (but all values can be treated as expressions $\langle v \rangle$). The partition results in some tedious rules in the reduction relation. For instance, $\langle v_1 \rangle ,, \langle v_2 \rangle \hookrightarrow \langle v_1 ,, v_2 \rangle$ reduces a merge of two values to a merged value.

6.2   Improvements and Extensions

Less Checks on Reduction.   In rule TREDUCE-ARROW (in Figure 4), the premise $C <: A$ is actually redundant for the purposes of reduction. Since we only care about well-typed terms being reduced, such a check has already been guaranteed by typing. Therefore an actual implementation could omit that check. The reason why we keep the premise is that typed reduction plays another role in our metatheory: it allows us to define consistency. Consistency is defined for any (untyped) values, and the extra check there tightens up the definition of consistency. With the premise, typed reduction directly implies a subtyping relation between the principal type of the reduced value and the reduction type. (See Lemma 9: If $v \hookrightarrow_A v'$, then $type_p\langle v \rangle <: A$). One could wonder if this property is unnecessary because it may be

derived by type preservation of reduction. Note that whenever typed reduction is called in a reduction rule, the subtyping relation can be obtained from the typing derivation of the reduced term. For example, reducing $v : A$ will type reduce v under A. If $v : A$ is well-typed, then we could in principle prove that $type_p \langle v \rangle <: A$. Unfortunately, the above proof is hard to attain in practice. Because type preservation depends on consistency, and consistency is defined by typed reduction. Once the subtyping property relies on type preservation, there is a cyclic dependency between the properties. In future work we would like to look at this issue more closely and try to discard the premise by taking full advantage of the type system.

Distributive Subtyping. Although the subtyping of $\lambda_i^:$ allows multiple functions in a merge, it lacks the distributive subtyping rule for intersection types that has been employed in some recent calculi [5,6]. The distributivity of intersections over arrows $((A \to B_1) \,\&\, (A \to B_2) <: A \to B_1 \,\&\, B_2)$ [3] is well accepted for its theoretical elegance. But it is also well-known for being troublesome. Mainly, there are two challenges for adapting distributive subtyping to $\lambda_i^:$.

- The rule indicates that a merge of functions can be applied. While the current typing rule can check such application with suitable annotations, designing new reduction rules is necessary. A promising solution is to have a rule allows parallel application like $(v_1 \,,, v_2) \, v \; \hookrightarrow \; v_1 \, v \,,, v_2 \, v$.
- Function types are no longer "ordinary". In $\lambda_i^:$, the intuition behind ordinary types is that their typed reduction results never contains merges, which is necessary for determinism. With distributivity, typed reduction may produce a merge under a single function type. For example, $\lambda x. \,'c' : \mathsf{Int} \to \mathsf{Char} \,,, \lambda x. \, x : \mathsf{Int} \to \mathsf{Int} \; \hookrightarrow_{\mathsf{Int} \to \mathsf{Char} \,\&\, \mathsf{Int}} \; \lambda x. \,'c' : \mathsf{Int} \to \mathsf{Char} \,,, \lambda x. \, x : \mathsf{Int} \to \mathsf{Int}$. In the typed reduction of $\lambda_i^:$, intersections are split into basic units. However, it is not straightforward to split a function type.

7 Related Work

7.1 Calculi with the Merge Operator and a Direct Semantics

Intersection types with a merge operator are a key feature of Reynolds' Forsythe language [43]. Reynolds studied a core calculus [43] with similarities to $\lambda_i^:$. However, merges in Forsythe are restricted and use a syntactic criterion to determine what merges are allowed. A merge is permitted only when the second term is a lambda abstraction or a single field record, which makes the structure of merge always biased. To prevent potential ambiguity, the latter overrides the former when overlapped. Note that the structure of merge in Forsythe is always biased. If formalized as a tree, the right child of every node is a leaf. The only place for primitive types is the leftmost component. Forsythe follows the standard call-by-name small-step reduction, during which types are ignored. The reduction rules deal with merges by continuously checking if the second component can be used in the context (abstractions for application, records for projection). This simple approach, however, is unable to reduce merges when (multiple) primitive types are required. Reynolds admits this issue in his later work [45]. In $\lambda_i^:$ types are used to select values from a merge and the disjointness restriction guarantees the determinism. Therefore the order of a value in a merge is not a deciding factor on whether the value is used or not.

The calculus $\lambda\&$ proposed by Castagna et al. [9] has a restricted version of the merge operator for functions only. The merge operator is indexed by a list of types of its components. The operational semantics uses the run-time types of values to select the "best approximate" branch of an overloaded function. $\lambda\&$ requires run-time type checking on values, while

	Dunfield's [23]	λ_i [20]	F_i [2]	NeColus [5]	F_i^+ [6]	$\lambda_i^:$
Disjointness	○	●	●	●	●	●
Unrestricted Intersections	●	○	○	●	●	●
Determinism or Coherence	No	Coh.	Coh.	Coh.	Coh.	Det.
Coercion Free	●	○	○	○	○	●
Recursion	●	○	○	○	○	●
Direct Semantics	●	○	○	○	○	●
Subject-Reduction	○	-	-	-	-	●

Figure 8 Summary of intersection calculi with the merge operator (● = yes, ○ = no, - = not applicable).

in TDOS, all type information is present already in type annotations. Another obvious difference is that $\lambda_i^:$ supports merges of any types (not just functions), which are useful for applications other than overloading of functions, including: multifield *extensible records with subtyping* [20]; encodings of *objects* and *traits* [4]; *dynamic mixins* [2]; or simple forms of *family polymorphism* [5].

Several other calculi with intersection types and overloading of functions have been proposed [10–12], but these calculi do not support a merge operator, and thus avoid the ambiguity problems caused by the construct.

7.2   Calculi with a Merge Operator and an Elaboration Semantics

Instead of a direct semantics, many recent works [2, 5, 6, 20, 23] on intersection types employ an elaboration semantics, translating merges in the source language to products (or pairs) in a target language. With an elaboration semantics the subtyping derivations are coercive [33]: they produce coercion functions that explicitly convert terms of one type to another in the target language. This idea was first proposed by Dunfield [23], where she shows how to elaborate a calculus with intersection and union types and a merge operator to a standard call-by-value lambda calculus with products and sums. Dunfield also proposed a direct semantics, which served as inspiration for our own work. However, her direct semantics is non-deterministic and lacks subject-reduction (as discussed in detail in Section 2.2). Unlike Forsythe and λ&, Dunfield's calculus has unrestricted merges and allows a merge to work as an argument. Her calculus is flexible and expressive and can deal with several programs that are not allowed in Forsythe and λ&.

To remove the ambiguity issues in Dunfield's work, the λ_i calculus [20] forbids overlapping in intersections using the disjointness restriction for all well-formed intersections. In other words, λ_i does not support unrestricted intersections. Because of this restriction, the proof of coherence in λ_i is still relatively simple. Likewise, in following work on the F_i calculus [2], which extends λ_i with disjoint polymorphism, *all* intersections must be *disjoint*. However the disjointness restriction causes difficulties because it breaks *stability of type substitutions*. Stability is a desirable property in a polymorphic type system that ensures that if a polymorphic type is well-formed then any instantiation of that type is also well-formed. Unfortunately, with disjoint intersections only, this property is not true in general. Thus F_i can only prove a restricted version of stability, which makes its metatheory non-trivial.

Disjointness of all well-formed intersections is only a sufficient (but not necessary) restriction to ensure an unambiguous semantics. The NeColus calculus [5] relaxes the restriction without introducing ambiguity. In NeColus 1 : Int & Int is allowed, but the same term is

rejected in λ_i. In other words, NeColus employs the disjointness restriction *only* on merges, but otherwise allows *unrestricted intersections*. Unfortunately, this comes at a cost: it is much harder to prove the coherence of elaboration. Both NeColus and F_i^+ [6] (a calculus derived from F_i that allows unrestricted intersections) deal with this problem by establishing coherence using contextual equivalence and a *logical relation* [40, 49, 50] to prove it. The proof method, however, cannot deal with non-terminating programs. In fact none of the existing calculi with disjoint intersection types supports recursion, which is a severe restriction.

We retain the essence of the power of Dunfield's calculus (modulo the disjointness restrictions to rule out ambiguity), and gain benefits from the direct semantics. Figure 8 summarizes the key differences between our work and prior work, focusing on the most recent work on disjoint intersection types. Note that the row titled "Coercion Free" denotes whether subtyping generates coercions or not. $\lambda_i^{\;:}$ is coercion free, while all other calculi based on an elaboration semantics employ coercive subtyping. Next we give more detail on the advantages of a direct semantics over the elaboration semantics and proof methods employed in previous work on disjoint intersection types.

Shorter, more Direct Reasoning. Programmers want to understand the meaning of their programs. A formal semantics can help with this. With our TDOS semantics we can essentially employ a style similar to equational reasoning in functional programming to directly reason about programs written in $\lambda_i^{\;:}$. For example, it takes a few reasoning steps to work out the result of $(\lambda x.\, x\, +\, 1 : \mathsf{Int} \to \mathsf{Int})\, (2\, ,, \, \text{‘}c\text{’})$:

$\quad (\lambda x.\, x\, +\, 1 : \mathsf{Int} \to \mathsf{Int})\, (2\, ,, \, \text{‘}c\text{’})$
$\hookrightarrow (2 + 1) : \mathsf{Int}$ {by STEP-BETA and typed reduction of argument under Int}
$\hookrightarrow 3 : \mathsf{Int}$ {by STEP-ANNO and usual reduction rules for arithmetic}
$\hookrightarrow 3$ {by STEP-ANNOV and typed reduction of 3 under Int}

Here reasoning is easily justifiable from the small-step reduction rules and type-directed reduction. In fact building tools (such as some form of debugger), that automate such kind of reasoning should be easy using the TDOS rules.

However, with an elaboration semantics, the (precise) reasoning steps to determine the final result are much more complex. Firstly the expression has to be translated into the target language before reducing to a similar target term. Figure 9 shows this elaboration process in λ_i, where an expression in the source language is translated into an expression in a target language with products. The source term $(\lambda x.\, x\, +\, 1 : \mathsf{Int} \to \mathsf{Int})\, (2\, ,, \, \text{‘}c\text{’})$ is elaborated into the target term $(\lambda x.\, x\, +\, 1)\, (\mathsf{fst}\, (2\, , \text{‘}c\text{’}))$. As we can see the actual derivation is rather long, so we skip the full steps. Also, for simplicity's sake, here we assume the subtyping judgement produces the most straightforward coercion fst. This elaboration step together with the introduction of coercions into the program makes it much harder for programmers to precisely understand the semantics of a program. Moreover while the coercions inserted in this small expression may not look too bad, in larger programs the addition of coercions can be a lot more severe, hampering the understanding of the program. After elaboration we can then use the target language semantics, to determine a target language value.

$\quad (\lambda x.\, x\, +\, 1)\, (\mathsf{fst}\, (2\, , \text{‘}c\text{’}))$
$\hookrightarrow (\lambda x.\, x\, +\, 1)\, 2$ {by a rule similar to STEP-APPR and reduction rules for pairs}
$\hookrightarrow 2 + 1$ {by beta reduction rule}
$\hookrightarrow 3$ {by usual reduction rules for arithmetic}

$$
\begin{array}{c}
\text{T-MERGE} \\
\cfrac{\cdots}{\cdot \vdash (2\,,\,\text{`}c\text{'}) \Rightarrow \mathsf{Int}\,\&\,\mathsf{Char} \rightsquigarrow (2\,,\text{`}c\text{'})}
\end{array}
$$

$$
\text{T-ANN} \quad
\cfrac{\cdots}{\cdot \vdash (\lambda x.\,x+1 : \mathsf{Int} \to \mathsf{Int}) \Rightarrow \mathsf{Int} \to \mathsf{Int} \rightsquigarrow \lambda x.\,x+1}
$$

$$
\text{SUB-ANDL} \quad \mathsf{Int}\,\&\,\mathsf{Char} <: \mathsf{Int} \rightsquigarrow \mathsf{fst}
$$

$$
\text{T-SUB} \;\; \cfrac{}{\cdot \vdash (2\,,\,\text{`}c\text{'}) \Leftarrow \mathsf{Int} \rightsquigarrow \mathsf{fst}\,(2\,,\text{`}c\text{'})}
$$

$$
\text{T-APP} \;\; \cfrac{}{\cdot \vdash (\lambda x.\,x+1 : \mathsf{Int} \to \mathsf{Int})\,(2\,,\,\text{`}c\text{'}) \Rightarrow \mathsf{Int} \rightsquigarrow (\lambda x.\,x+1)\,(\mathsf{fst}\,(2\,,\text{`}c\text{'}))}
$$

■ **Figure 9** Elaboration of the expression $(\lambda x.\,x+1 : \mathsf{Int} \to \mathsf{Int})\,(2\,,\,\text{`}c\text{'})$ into a target calculus with products.

A final issue is that sometimes it is not even possible to translate back the value of the target language into an equivalent "value" on the source. For instance in the NeColus calculus [5] $1 : \mathsf{Int}\,\&\,\mathsf{Int}$ results in $(1,\,1)$, which is a pair in the target language. But the corresponding source value $1\,,\,1$ is not typable in NeColus. In essence, with an elaboration, programmers must understand not only the source language, but also the elaboration process as well as the semantics of the target language, if they want to precisely understand the semantics of a program. Since the main point of semantics is to give clear and simple rules to understand the meaning of programs, a direct semantics is a better option for providing such understanding.

Simpler Proofs of Unambiguity.   For calculi with an elaboration semantics, unrestricted intersections make it harder to prove the coherence. Our $\lambda_i^{:}$ calculus, on the other hand, has a deterministic semantics, which implies unambiguity directly. For instance, $(1 : \mathsf{Int}\,\&\,\mathsf{Int}) :$ Int only steps to 1 in $\lambda_i^{:}$. But it can be elaborated into two target expressions in the NeColus calculus corresponding to two typing derivations:

$$(1 : \mathsf{Int}\,\&\,\mathsf{Int}) : \mathsf{Int} \rightsquigarrow \mathsf{fst}\,(1,\,1)$$

$$(1 : \mathsf{Int}\,\&\,\mathsf{Int}) : \mathsf{Int} \rightsquigarrow \mathsf{snd}\,(1,\,1)$$

Thus the coherence proof needs deeper knowledge about the semantics: the two different terms are known to both reduce to 1 eventually. Therefore they are related by the logical relation employed in NeColus for coherence. Things get more complicated for functions. The following example shows two possible elaborations of the same function. To relate them requires reasoning inside the binders and a notion of contextual equivalence.

$$\lambda x.\,x+1 : \mathsf{Int}\,\&\,\mathsf{Int} \to \mathsf{Int} \rightsquigarrow \lambda x.\,\mathsf{fst}\;x+1$$

$$\lambda x.\,x+1 : \mathsf{Int}\,\&\,\mathsf{Int} \to \mathsf{Int} \rightsquigarrow \lambda x.\,\mathsf{snd}\;x+1$$

Furthermore, the two target expressions above are *clearly not equivalent* in the general case. For instance, if we apply them to $(1,\,2)$ we get different results. However, the target expressions will always behave equivalently when applied to arguments *elaborated from the* NeColus *source calculus*. NeColus, forbids terms like $(1\,,\,2)$ and thus cannot produce a target value $(1,\,2)$. Because of elaboration and also this deeper form of reasoning required to show the equivalence of semantics, calculi defined by elaboration require a lot more infrastructure for the source and target calculi and the elaboration between them, while in a direct semantics only one calculus is involved and the reasoning required to prove determinism is quite simple.

Not Limited to Terminating Programs. The (basic) forms of logical relations employed by NeColus and F_i^+ has cannot deal with non-terminating programs. In principle, recursion could be supported by using a *step-indexed logical relation* [1], but this is left for future work (and it is non-trivial). $\lambda_i^{:}$ smoothly handles unrestricted intersections and recursion, using TDOS to reach determinism with a significantly simpler proof method. It also makes other features that lead to non-terminating programs, such as *recursive types*, feasible.

7.3 Languages and Calculi with Type-Dependent Semantics

Typed Operational Semantics Goguen [29] uses types in its reduction judgment, similarly to typed reduction in $\lambda_i^{:}$. However, Goguen's typed operational semantics is designed for studying meta-theoretic properties, especially strong normalization, and is not aimed to describe type-dependent semantics. Unlike TDOS, in typed operational semantics the reduction process does not use the additional type information to guide reduction. Instead, the combination of well-typedness and computation provides inversion principles for proving various metatheoretical properties. Typed operational semantics has been applied to several systems. These include *simply typed lambda calculi* [30], calculi with *dependent types* [26, 29] and *higher-order subtyping* [17]. Note that the semantics of these systems does not depend on typing, and the untyped (type-erased) reduction relations are still presented to describe how to evaluate programs.

Type classes [32, 52] are an approach to parametric overloading used in languages like Haskell. The commonly adopted compilation strategy for it is the dictionary passing style elaboration [13, 14, 31, 52]. Other mechanisms inspired by type classes, such as Scala's *implicits* [19], Agda's *instance arguments* [22] or Ocaml's *modular implicits* [54] have an elaboration semantics as well. In one of the pioneering works of type classes, Kaes [32] gives two formulations for a direct operational semantics. One of them decides the concrete type of the instance of overloaded functions at *run-time*, by analyzing all arguments after evaluating them. In both Kaes' work and a following work by Odersky et al. [37], the run-time semantics has some restrictions with respect to type classes. For example, overloading on return types (needed for example for the *read* function in Haskell) is not supported. Interestingly, the semantics of $\lambda_i^{:}$ allows overloading on return types, which is used whenever two functions coexist on a merge.

Gradual typing [47] has become popular over the last few years. Gradual typing is another example of a type-dependent mechanism, since the success or not of an (implicit) cast may depend on the particular type used for the implicit cast. Thus the semantics of a gradually typed language is type-dependent. Like other type-dependent mechanisms the semantics of gradually typed source languages is usually given by a (type-dependent) elaboration semantics into a cast calculus, such as the *Blame calculus* [53] or the *Threesome calculus* [48].

Static binding of fields and *method overloading* in Java [51] make use of type annotations computed in a preprocessing phase. For each method invocation, the annotation states the argument type of the most specific method applicable according to the static types. Based on the annotation and the run-time type (class) of the object, a dynamic lookup function yields a proper method at run-time. This allows static method overloading works across the inheritance hierarchy, together with dynamic dispatch. *Multiple dispatching* [15, 16, 35, 38] generalizes object-oriented dynamic dispatch to determine the overloaded method to invoke based on the run-time type of all its arguments. Similarly to TDOS, much of the type information is recovered from type annotations in multiple dispatching mechanisms, but, unlike TDOS, they only use input types to determine the semantics.

8   Conclusion

In this work we presented a TDOS for $\lambda_i^{:}$: a calculus that includes intersection types and an expressive unbiased merge operator. Among all similar calculi, $\lambda_i^{:}$ is the first to have a direct operational semantics that is both deterministic and has subject-reduction. Compared with the elaboration approach, having a direct semantics avoids the translation process and a target calculus. This simplifies both informal and formal reasoning. For instance, establishing the coherence of elaboration in NeColus [5] requires much more sophistication than obtaining the determinism theorem in $\lambda_i^{:}$. Furthermore the proof method for coherence in NeColus cannot deal with non-terminating programs, whereas dealing with recursion is straightforward in $\lambda_i^{:}$. The semantics of $\lambda_i^{:}$ exploits type annotations to guide reduction. The key component of TDOS is *typed reduction*, which allows values to be further reduced depending on their type. For the future we would like to develop further the TDOS approach in the setting of disjoint intersection types. Some interesting extensions include support for *distributive subtyping* [3], *disjoint polymorphism* [2] and iso-recursive types with the Amber rule [7].

References

1   Amal J. Ahmed. Step-indexed syntactic logical relations for recursive and quantified types. In Peter Sestoft, editor, *Programming Languages and Systems, 15th European Symposium on Programming, ESOP 2006, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2006, Vienna, Austria, March 27-28, 2006, Proceedings*, volume 3924 of *Lecture Notes in Computer Science*, pages 69–83. Springer, 2006. `doi: 10.1007/11693024_6`.

2   João Alpuim, Bruno C. d. S. Oliveira, and Zhiyuan Shi. Disjoint polymorphism. In Hongseok Yang, editor, *Programming Languages and Systems - 26th European Symposium on Programming, ESOP 2017, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2017, Uppsala, Sweden, April 22-29, 2017, Proceedings*, volume 10201 of *Lecture Notes in Computer Science*, pages 1–28. Springer, 2017. `doi:10.1007/978-3-662-54434-1_1`.

3   Henk Barendregt, Mario Coppo, and Mariangiola Dezani-Ciancaglini. A filter lambda model and the completeness of type assignment 1. *The journal of symbolic logic*, 48(4), 1983.

4   Xuan Bi and Bruno C. d. S. Oliveira. Typed first-class traits. In Todd D. Millstein, editor, *32nd European Conference on Object-Oriented Programming, ECOOP 2018, July 16-21, 2018, Amsterdam, The Netherlands*, volume 109 of *LIPIcs*, pages 9:1–9:28. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018. `doi:10.4230/LIPIcs.ECOOP.2018.9`.

5   Xuan Bi, Bruno C. d. S. Oliveira, and Tom Schrijvers. The essence of nested composition. In Todd D. Millstein, editor, *32nd European Conference on Object-Oriented Programming, ECOOP 2018, July 16-21, 2018, Amsterdam, The Netherlands*, volume 109 of *LIPIcs*, pages 22:1–22:33. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018. `doi:10.4230/LIPIcs. ECOOP.2018.22`.

6   Xuan Bi, Ningning Xie, Bruno C. d. S. Oliveira, and Tom Schrijvers. Distributive disjoint polymorphism for compositional programming. In Luís Caires, editor, *Programming Languages and Systems - 28th European Symposium on Programming, ESOP 2019, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2019, Prague, Czech Republic, April 6-11, 2019, Proceedings*, volume 11423 of *Lecture Notes in Computer Science*, pages 381–409. Springer, 2019. `doi:10.1007/978-3-030-17184-1_14`.

7   Luca Cardelli. Amber. In Guy Cousineau, Pierre-Louis Curien, and Bernard Robinet, editors, *Combinators and Functional Programming Languages, Thirteenth Spring School of the LITP, Val d'Ajol, France, May 6-10, 1985, Proceedings*, volume 242 of *Lecture Notes in Computer Science*, pages 21–47. Springer, 1985. `doi:10.1007/3-540-17184-3_38`.

8   Luca Cardelli and Peter Wegner. On understanding types, data abstraction, and polymorphism. *ACM Comput. Surv.*, 17(4):471–522, 1985. `doi:10.1145/6041.6042`.

9   Giuseppe Castagna, Giorgio Ghelli, and Giuseppe Longo. A calculus for overloaded functions with subtyping. *Inf. Comput.*, 117(1):115–135, 1995. `doi:10.1006/inco.1995.1033`.

10 Giuseppe Castagna, Kim Nguyen, Zhiwu Xu, and Pietro Abate. Polymorphic functions with set-theoretic types: Part 2: Local type inference and type reconstruction. In Sriram K. Rajamani and David Walker, editors, *Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2015, Mumbai, India, January 15-17, 2015*, pages 289–302. ACM, 2015. `doi:10.1145/2676726.2676991`.

11 Giuseppe Castagna, Kim Nguyen, Zhiwu Xu, Hyeonseung Im, Sergueï Lenglet, and Luca Padovani. Polymorphic functions with set-theoretic types: part 1: syntax, semantics, and evaluation. In Suresh Jagannathan and Peter Sewell, editors, *The 41st Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '14, San Diego, CA, USA, January 20-21, 2014*, pages 5–18. ACM, 2014. `doi:10.1145/2535838.2535840`.

12 Giuseppe Castagna and Zhiwu Xu. Set-theoretic foundation of parametric polymorphism and subtyping. In Manuel M. T. Chakravarty, Zhenjiang Hu, and Olivier Danvy, editors, *Proceeding of the 16th ACM SIGPLAN international conference on Functional Programming, ICFP 2011, Tokyo, Japan, September 19-21, 2011*, pages 94–106. ACM, 2011. `doi:10.1145/2034773.2034788`.

13 Manuel M. T. Chakravarty, Gabriele Keller, and Simon L. Peyton Jones. Associated type synonyms. In Olivier Danvy and Benjamin C. Pierce, editors, *Proceedings of the 10th ACM SIGPLAN International Conference on Functional Programming, ICFP 2005, Tallinn, Estonia, September 26-28, 2005*, pages 241–253. ACM, 2005. `doi:10.1145/1086365.1086397`.

14 Manuel M. T. Chakravarty, Gabriele Keller, Simon L. Peyton Jones, and Simon Marlow. Associated types with class. In Jens Palsberg and Martín Abadi, editors, *Proceedings of the 32nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2005, Long Beach, California, USA, January 12-14, 2005*, pages 1–13. ACM, 2005. `doi:10.1145/1040305.1040306`.

15 Craig Chambers and Weimin Chen. Efficient multiple and predicated dispatching. In Brent Hailpern, Linda M. Northrop, and A. Michael Berman, editors, *Proceedings of the 1999 ACM SIGPLAN Conference on Object-Oriented Programming Systems, Languages & Applications (OOPSLA '99), Denver, Colorado, USA, November 1-5, 1999*, pages 238–255. ACM, 1999. `doi:10.1145/320384.320407`.

16 Curtis Clifton, Gary T. Leavens, Craig Chambers, and Todd D. Millstein. Multijava: modular open classes and symmetric multiple dispatch for java. In Mary Beth Rosson and Doug Lea, editors, *Proceedings of the 2000 ACM SIGPLAN Conference on Object-Oriented Programming Systems, Languages & Applications (OOPSLA 2000), Minneapolis, Minnesota, USA, October 15-19, 2000*, pages 130–145. ACM, 2000. `doi:10.1145/353171.353181`.

17 Adriana B. Compagnoni and Healfdene Goguen. Typed operational semantics for higher-order subtyping. *Inf. Comput.*, 184(2):242–297, 2003. `doi:10.1016/S0890-5401(03)00062-2`.

18 Mario Coppo, Mariangiola Dezani-Ciancaglini, and Betti Venneri. Functional characters of solvable terms. *Math. Log. Q.*, 27(2-6):45–58, 1981. `doi:10.1002/malq.19810270205`.

19 Bruno C. d. S. Oliveira, Adriaan Moors, and Martin Odersky. Type classes as objects and implicits. In William R. Cook, Siobhán Clarke, and Martin C. Rinard, editors, *Proceedings of the 25th Annual ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications, OOPSLA 2010, October 17-21, 2010, Reno/Tahoe, Nevada, USA*, pages 341–360. ACM, 2010. `doi:10.1145/1869459.1869489`.

20 Bruno C. d. S. Oliveira, Zhiyuan Shi, and João Alpuim. Disjoint intersection types. In Jacques Garrigue, Gabriele Keller, and Eijiro Sumii, editors, *Proceedings of the 21st ACM SIGPLAN International Conference on Functional Programming, ICFP 2016, Nara, Japan, September 18-22, 2016*, pages 364–377. ACM, 2016. `doi:10.1145/2951913.2951945`.

21   Rowan Davies and Frank Pfenning. Intersection types and computational effects. In Martin Odersky and Philip Wadler, editors, *Proceedings of the Fifth ACM SIGPLAN International Conference on Functional Programming (ICFP '00), Montreal, Canada, September 18-21, 2000*, pages 198–208. ACM, 2000. `doi:10.1145/351240.351259`.

22   Dominique Devriese and Frank Piessens. On the bright side of type classes: instance arguments in agda. In Manuel M. T. Chakravarty, Zhenjiang Hu, and Olivier Danvy, editors, *Proceeding of the 16th ACM SIGPLAN international conference on Functional Programming, ICFP 2011, Tokyo, Japan, September 19-21, 2011*, pages 143–155. ACM, 2011. `doi:10.1145/2034773.2034796`.

23   Jana Dunfield. Elaborating intersection and union types. *J. Funct. Program.*, 24(2-3):133–165, 2014. `doi:10.1017/S0956796813000270`.

24   Jana Dunfield and Frank Pfenning. Type assignment for intersections and unions in call-by-value languages. In Andrew D. Gordon, editor, *Foundations of Software Science and Computational Structures, 6th International Conference, FOSSACS 2003 Held as Part of the Joint European Conference on Theory and Practice of Software, ETAPS 2003, Warsaw, Poland, April 7-11, 2003, Proceedings*, volume 2620 of *Lecture Notes in Computer Science*, pages 250–266. Springer, 2003. `doi:10.1007/3-540-36576-1_16`.

25   Facebook. Flow. `https://flow.org/`, 2014.

26   Yangyue Feng and Zhaohui Luo. Typed operational semantics for dependent record types. In Tom Hirschowitz, editor, *Proceedings Types for Proofs and Programs, Revised Selected Papers, TYPES 2009, Aussois, France, 12-15th May 2009*, volume 53 of *EPTCS*, pages 30–46, 2009. `doi:10.4204/EPTCS.53.3`.

27   Matthew Flatt, Shriram Krishnamurthi, and Matthias Felleisen. Classes and mixins. In David B. MacQueen and Luca Cardelli, editors, *POPL '98, Proceedings of the 25th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, San Diego, CA, USA, January 19-21, 1998*, pages 171–183. ACM, 1998. `doi:10.1145/268946.268961`.

28   Timothy S. Freeman and Frank Pfenning. Refinement types for ML. In David S. Wise, editor, *Proceedings of the ACM SIGPLAN'91 Conference on Programming Language Design and Implementation (PLDI), Toronto, Ontario, Canada, June 26-28, 1991*, pages 268–277. ACM, 1991. `doi:10.1145/113445.113468`.

29   Healfdene Goguen. *A typed operational semantics for type theory*. PhD thesis, University of Edinburgh, UK, 1994. URL: `http://hdl.handle.net/1842/405`.

30   Healfdene Goguen. Typed operational semantics. In Mariangiola Dezani-Ciancaglini and Gordon D. Plotkin, editors, *Typed Lambda Calculi and Applications, Second International Conference on Typed Lambda Calculi and Applications, TLCA '95, Edinburgh, UK, April 10-12, 1995, Proceedings*, volume 902 of *Lecture Notes in Computer Science*, pages 186–200. Springer, 1995. `doi:10.1007/BFb0014053`.

31   Cordelia V. Hall, Kevin Hammond, Simon L. Peyton Jones, and Philip Wadler. Type classes in haskell. *ACM Trans. Program. Lang. Syst.*, 18(2):109–138, 1996. `doi:10.1145/227699.227700`.

32   Stefan Kaes. Parametric overloading in polymorphic programming languages. In Harald Ganzinger, editor, *ESOP '88, 2nd European Symposium on Programming, Nancy, France, March 21-24, 1988, Proceedings*, volume 300 of *Lecture Notes in Computer Science*, pages 131–144. Springer, 1988. `doi:10.1007/3-540-19027-9_9`.

33   Zhaohui Luo. Coercive subtyping. *J. Log. Comput.*, 9(1):105–130, 1999. `doi:10.1093/logcom/9.1.105`.

34   Microsoft. Typescript. `https://www.typescriptlang.org/`, 2012.

35   Radu Muschevici, Alex Potanin, Ewan D. Tempero, and James Noble. Multiple dispatch in practice. In Gail E. Harris, editor, *Proceedings of the 23rd Annual ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications, OOPSLA 2008, October 19-23, 2008, Nashville, TN, USA*, pages 563–582. ACM, 2008. `doi:10.1145/1449764.1449808`.

36   Martin Odersky, Philippe Altherr, Vincent Cremet, Burak Emir, Sebastian Maneth, Stéphane Micheloud, Nikolay Mihaylov, Michel Schinz, Erik Stenman, and Matthias Zenger. An overview of the scala programming language. Technical report, École Polytechnique Fédérale de Lausanne, 2004.

37   Martin Odersky, Philip Wadler, and Martin Wehr. A second look at overloading. In John Williams, editor, *Proceedings of the seventh international conference on Functional programming languages and computer architecture, FPCA 1995, La Jolla, California, USA, June 25-28, 1995*, pages 135–146. ACM, 1995. `doi:10.1145/224164.224195`.

38   Gyunghee Park, Jaemin Hong, Guy L. Steele Jr., and Sukyoung Ryu. Polymorphic symmetric multiple dispatch with variance. *Proc. ACM Program. Lang.*, 3(POPL):11:1–11:28, 2019. `doi:10.1145/3290324`.

39   Benjamin C. Pierce. *Programming with Intersection Types and Bounded Polymorphism*. PhD thesis, Carnegie Mellon University, December 1991.

40   Gordon Plotkin. Lambda-definability and logical relations, 1973.

41   Garrel Pottinger. A type assignment for the strongly normalizable λ-terms. *To HB Curry: essays on combinatory logic, lambda calculus and formalism*, pages 561–577, 1980.

42   Redhat. Ceylon. `https://ceylon-lang.org/`, 2011.

43   John C Reynolds. Preliminary design of the programming language Forsythe. Technical Report CMU-CS-88-159, Carnegie Mellon University, 1988.

44   John C. Reynolds. The coherence of languages with intersection types. In Takayasu Ito and Albert R. Meyer, editors, *Theoretical Aspects of Computer Software, International Conference TACS '91, Sendai, Japan, September 24-27, 1991, Proceedings*, volume 526 of *Lecture Notes in Computer Science*, pages 675–700. Springer, 1991. `doi:10.1007/3-540-54415-1_70`.

45   John C Reynolds. Design of the programming language Forsythe. In *ALGOL-like languages*, pages 173–233. Springer, 1997.

46   Nathanael Schärli, Stéphane Ducasse, Oscar Nierstrasz, and Andrew P. Black. Traits: Composable units of behaviour. In Luca Cardelli, editor, *ECOOP 2003 - Object-Oriented Programming, 17th European Conference, Darmstadt, Germany, July 21-25, 2003, Proceedings*, volume 2743 of *Lecture Notes in Computer Science*, pages 248–274. Springer, 2003. `doi:10.1007/978-3-540-45070-2_12`.

47   Jeremy G Siek and Walid Taha. Gradual typing for functional languages. In *Scheme and Functional Programming Workshop*, 2006.

48   Jeremy G. Siek and Philip Wadler. Threesomes, with and without blame. In Manuel V. Hermenegildo and Jens Palsberg, editors, *Proceedings of the 37th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2010, Madrid, Spain, January 17-23, 2010*, pages 365–376. ACM, 2010. `doi:10.1145/1706299.1706342`.

49   Richard Statman. Logical relations and the typed λ-calculus. *Inf. Control.*, 65(2/3):85–97, 1985. `doi:10.1016/S0019-9958(85)80001-2`.

50   William W. Tait. Intensional interpretations of functionals of finite type I. *J. Symb. Log.*, 32(2):198–212, 1967. `doi:10.2307/2271658`.

51   David von Oheimb and Tobias Nipkow. Machine-checking the java specification: Proving type-safety. In Jim Alves-Foss, editor, *Formal Syntax and Semantics of Java*, volume 1523 of *Lecture Notes in Computer Science*, pages 119–156. Springer, 1999. `doi:10.1007/3-540-48737-9_4`.

52   Philip Wadler and Stephen Blott. How to make ad-hoc polymorphism less ad-hoc. In *Conference Record of the Sixteenth Annual ACM Symposium on Principles of Programming Languages, Austin, Texas, USA, January 11-13, 1989*, pages 60–76. ACM Press, 1989. `doi:10.1145/75277.75283`.

53   Philip Wadler and Robert Bruce Findler. Well-typed programs can't be blamed. In Giuseppe Castagna, editor, *Programming Languages and Systems, 18th European Symposium on Programming, ESOP 2009, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2009, York, UK, March 22-29, 2009. Proceedings*, volume 5502 of *Lecture Notes in Computer Science*, pages 1–16. Springer, 2009. `doi:10.1007/978-3-642-00590-9_1`.

54   Leo White, Frédéric Bour, and Jeremy Yallop. Modular implicits. In Oleg Kiselyov and Jacques Garrigue, editors, *Proceedings ML Family/OCaml Users and Developers workshops, ML/OCaml 2014, Gothenburg, Sweden, September 4-5, 2014*, volume 198 of *EPTCS*, pages 22–63, 2014. `doi:10.4204/EPTCS.198.2`.

55   Andrew K. Wright and Matthias Felleisen. A syntactic approach to type soundness. *Inf. Comput.*, 115(1):38–94, 1994. `doi:10.1006/inco.1994.1093`.