A Bowtie for a Beast

Overloading, Eta Expansion, and Extensible Data Types in F⋈

NICK RIOUX, University of Pennsylvania, USA
XUEJING HUANG and BRUNO C. D. S. OLIVEIRA, The University of Hong Kong, China
STEVE ZDANCEWIC, University of Pennsylvania, USA

The typed *merge operator* offers the promise of a compositional style of statically-typed programming in which solutions to the expression problem arise naturally. This approach, dubbed *compositional programming*, has recently been demonstrated by Zhang et al.

Unfortunately, the merge operator is an unwieldy beast. Merging values from overlapping types may be ambiguous, so *disjointness relations* have been introduced to rule out undesired nondeterminism and obtain a well-behaved semantics. Past type systems using a disjoint merge operator rely on intersection types, but extending such systems to include union types or overloaded functions is problematic: naively adding either reintroduces ambiguity. In a nutshell: the elimination forms of unions and overloaded functions require values to be distinguishable by case analysis, but the merge operator can create exotic values that violate that requirement.

This paper presents F⋈, a core language that demonstrates how unions, intersections, and overloading can all coexist with a tame merge operator. The key is an underlying design principle that states that any two inhabited types can support either the deterministic merging of their values, or the ability to distinguish their values, but never both. To realize this invariant, we decompose previously studied notions of disjointness into two new, dual relations that permit the operation that best suits each pair of types. This duality respects the polarization of the type structure, yielding an expressive language that we prove to be both type safe and deterministic.

CCS Concepts: • **Software and its engineering** → **Functional languages**; **Data types and structures**; **Patterns**; • **Theory of computation** → **Type theory**.

Additional Key Words and Phrases: extensibility, polymorphism, type systems

1 INTRODUCTION

Given two programs e_1 and e_2, the merge operator $e_1 \ \parallel \ e_2$ combines them into one program that supports all the operations of both of its inputs.[1] This powerful language feature has been used to model a wide range of mechanisms. The canonical example is for record concatenation, which

---

[1]This operator is sometimes written $e_1 \ , , \ e_2$ in the literature, but we prefer the more symmetric "$\parallel$" since the operator is commutative and associative.

Authors' addresses: Nick Rioux, nrioux@cis.upenn.edu, University of Pennsylvania, Philadelphia, Pennsylvania, USA; Xuejing Huang, xjhuang@cs.hku.hk; Bruno C. d. S. Oliveira, bruno@cs.hku.hk, The University of Hong Kong, Hong Kong, China; Steve Zdancewic, stevez@cis.upenn.edu, University of Pennsylvania, Philadelphia, Pennsylvania, USA.

has long been used to model multiple inheritance [Cook et al. 1989; Wand 1989]. For instance, a multi-field record might be represented as the merge of three singleton records:

$$e = \{l_1 \mapsto 1\} \,\|\, \{l_2 \mapsto \mathbf{true}\} \,\|\, \{l_3 \mapsto 3\}$$

In a typed setting, one assigns merges intersection types. Thus, e has the type $\{l_1 : \mathbf{Int}\} \sqcap \{l_2 : \mathbf{Bool}\} \sqcap \{l_3 : \mathbf{Int}\}$. Here we are assuming the record fields are all distinct: $l_1 \neq l_2 \neq l_3$.

Unfortunately, the merge operator has a cost: merges of arbitrary values lead to nondeterminism. For example, while merging records with distinct fields, as in the program e above, is perfectly reasonable, merging the boolean values **true** and **false** would require the existence of a third boolean value that is both **true** and **false**. This violates programmers' expectation that only one branch of an if statement will execute: a program like **if true** $\|$ **false then** e_1 **else** e_2 is ambiguous. We would like a way to enable the power of the merge operator, while disallowing nondeterministic programs like this one.

Oliveira et al. [2016] introduce a typing discipline for the merge operator that rules out such nondeterministic programs using *disjoint intersection types*. Under such a typing discipline, only expressions with disjoint types can be merged. Since **Bool** is not disjoint from itself, the merge **true** $\|$ **false** is rejected. This line of work led to success in the modeling of a number of highly compositional programming patterns including forms of *family polymorphism* [Ernst 2001] and *first-class traits* [Bi and Oliveira 2018]. It has culminated in languages that support a style of programming known as *compositional programming* [Zhang et al. 2021] which offers solutions to challenges such as the Expression Problem [Wadler 1998]. At the heart of these designs is the ability of the merge operator to support *nested composition* [Bi et al. 2018], which exploits subtyping relationships like:

$$(\mathbf{Int} \rightarrow \{age : \mathbf{Int}\}) \sqcap (\mathbf{Int} \rightarrow \{name : \mathbf{String}\}) \quad \leq \quad \mathbf{Int} \rightarrow \{age : \mathbf{Int}, name : \mathbf{String}\}$$

Here, the merge operator lets us combine/compose the behavior of multiple implementations. If we think of the two functions in the intersection as constructors of objects, what one gets out is a new constructor for objects with the combined behaviors.

At the same time, the merge operator has also been shown to support a form of *function overloading*, in which a single function can have several bodies and the choice of which to run is determined by its arguments [Castagna et al. 1995; Dunfield 2014; Pierce 1991a]. The merge operator can combine multiple implementations of functions that take different argument types. Later, the appropriate implementation can be dispatched to based on the argument at the application site. However, overloading can also easily lead to nondeterminism if not carefully managed.

Both of these features, *nested composition* and *overloading*, are independently useful, but, so far, no system accommodates them simultaneously. The typing discipline used for the merge operator in Oliveira's line of work does not support traditional forms of overloading, and, conversely, the systems that support overloading do not support nested composition. This is the first problem that we address in this paper.

Another natural feature to consider alongside overloading is *union types* [Barbanera et al. 1995]. $A \sqcup B$ is the *untagged* union of A and B. Both A values and B values have this type (because, e.g., $A \leq A \sqcup B$), but, unlike for tagged unions, there is no extra dynamic information attached to the value to reveal which case it is. Nevertheless, the constructors of values of type A will often be distinct from those of type B, which is sufficient to tell them apart. For our purposes, we assume that the constructors of base types like **Int** are disjoint from those of other base types, like **Bool**, so 1 and **true** are distinct. In other words, a programmer can use a pattern-matching construct to eliminate a value e of type **Int** $\sqcup$ **Bool** as follows:

$$\mathbf{match}\; e \;\mathbf{with}\; \{(x_1 : \mathbf{Int}) \mapsto \mathbf{true}, (x_2 : \mathbf{Bool}) \mapsto \mathbf{false}\}$$

If $e = 1$: **Int**, then it is distinct from the constructors of **Bool**, so the match above will work as intended and evaluate to **true**. Likewise, if $e = $ **true** : **Bool**, the match will evaluate to **false**, so pattern matching discriminates between **Int** and **Bool** values as expected. Unfortunately, in the presence of the merge operator there is a problem: the term $e = 1 \,\|\,$ **true** has type **Int** $\sqcap$ **Bool**, which is a subtype of **Int** $\sqcup$ **Bool**. Now to decide whether e is an **Int** or a **Bool** during pattern matching yields the answer that it is *both*, leading again to nondeterministic evaluation. Safely incorporating union types with intersections and the merge operator is the second problem we address.

Disjointness and polarity. This paper gives a solution to both of the problems mentioned above. The key idea is that for certain types, like **Int** and **Bool** above, pattern matching provides a mechanism for distinguishing their values. When there is such a mechanism for types A and B, we say A and B are *distinguishable*, written $A \Leftrightarrow B$. Our language design requires that patterns in different branches of a match construct match distinguishable types. Types that are eliminated via pattern matching are naturally distinguishable from each other; such types are called *positive* in the literature on polarized type theory [Andreoli 1992]. Intuitively, the values of such types are defined completely by how they are constructed, and pattern matching extracts all of the information about them.

On the flip side, functions are characterized via a "strong" (i.e. non-matching) elimination form. They are characterized extensionally by the contexts in which they are used, not by how they are constructed. This means that function types are *negative*. Negative types are relevant because, to our knowledge, all practical uses of the merge operator in the literature concern merges of negative types. To apply the merge operator to two arguments of types A and B, our design requires that A and B are *mergeable*, written $A \bowtie B$.

Together, these two relations describe when it is unambiguous to merge two functions to create an overloaded function. For example, it makes sense for two function types such as **Int** $\rightarrow$ **String** and **Bool** $\rightarrow$ **String** to be mergeable because their inputs are distinguishable. Moreover, as typically (albeit not universally) described in lambda calculi, a function does not provide a means of dynamically asking whether it accepts integers or strings as inputs. Thus, from this perspective, we would *not* expect the type **Int** $\rightarrow$ **String** to be distinguishable from the type **Bool** $\rightarrow$ **String**, and in our system they are not distinguishable.

We refer to both distinguishability and mergeability as *disjointness relations*: the former is the disjointness of types as sets of values and the latter is the disjointness of types as sets of their contexts (i.e., the operations that can be carried out on their values).

A key invariant of our language design is that no two inhabited types are both distinguishable and mergeable—they can be only one or the other. This discipline rules out the issues of nondeterminism that arise in the presence of both intersection and union types. For instance, this principle rules out the ambiguous match for $e = 1 \,\|\,$ **true**: since **Int** and **Bool** are distinguishable, but they are not mergeable, and e is ill typed. (Indeed there is no term of the type **Int** $\sqcap$ **Bool** and we do not even consider that type to be well formed.) The observations that we have made about polarity suggest a natural answer to the question of whether a pair of types should be distinguishable (positive) or mergeable (negative). Our type system exploits this observation to create an expressive yet deterministic programming language.

Contributions. Our primary result is the formalization of a core language $F_{\bowtie}$ (pronounced "F bow"). We prove the type system sound and demonstrate its support for compositional programming. $F_{\bowtie}$ showcases several important aspects:

- It is the first language to include disjoint intersection and union types, overloading, and a deterministic merge operator. These features combine to permit uses of nested composition.

- It is also a step towards incorporating the work of Castagna, Ghelli, and Longo into the literature on disjoint intersection types [Oliveira et al. 2016].
- As in past treatments of the merge operator, the operational semantics is type-directed [Huang et al. 2021]. However, type information is *not* used for selecting which overload to execute at a call site. In $F_\bowtie$, dispatch is (co-)pattern matching. Consequently, we characterize the dynamic role of types as *runtime η-expansion*.
- The type system demonstrates the dual concepts of *mergeability* and *distinguishability*. Both have been studied independently [Oliveira et al. 2016; Rehman et al. 2022], but here it is shown that, in tandem, they enable well-typed deterministic overloading.

The key technical results in this paper have been proven using a combination of pencil-and-paper proofs for the results having to do with the term language constructs, and a Coq development. The paper proofs are available in this article's companion technical appendix [Rioux et al. 2022b]. The Coq development [Rioux et al. 2022a] formalizes certain type-level parts of the semantics including subtyping, dispatch, and some key properties of disjointness.

2 OVERVIEW

This section gives an overview of this paper, starting with background on merges and disjoint intersection types, and then introducing the key ideas of our work.

2.1 Background

Intersection and union types. Intersection and union types [Barbanera et al. 1995; Coppo and Dezani-Ciancaglini 1978; Pottinger 1980] are widely used in diverse fields of programming languages. Intersection types were introduced to characterize exactly all strongly normalizing lambda terms. Union types were later introduced as the dual construct of intersection types [Barbanera et al. 1995]. Intersection types were first adopted for programming by work on Forsythe [Reynolds 1988, 1997] and subsequently employed to express key aspects of multiple inheritance [Compagnoni and Pierce 1996] in object-oriented programming. The Scala language [Odersky et al. 2004] and its DOT calculus [Rompf and Amin 2016], for example, make fundamental use of intersection types to express a class/trait that extends multiple other traits. Union types have also been adopted in programming languages. For instance they are widely used in TypeScript and Flow, and were also included in Scala 3.

The merge operator. The Forsythe language [Reynolds 1988, 1997] introduced a *merge operator*, which allows building values that can have multiple types (expressed as intersection types). The merge operator has been studied more recently by Dunfield [2014], who removed significant restrictions originally present in Reynolds' design. A simple example of a program using the merge operator is:

$$\textbf{let } f = \textit{isDigit} \amalg \textit{not} \textbf{ in } (f \text{ '1'}, f \textbf{ false})$$

Here f is an overloaded function that can take either a character or a boolean as an argument; it has the type $(\textbf{Char} \rightarrow \textbf{Bool}) \sqcap (\textbf{Bool} \rightarrow \textbf{Bool})$. The variable f is built using the merge operator and applying it extracts one of the functions from the merged value. In the body of the **let**, we see two applications of f: one to a character, and another to a boolean. This style of overloading is one of the major features of the merge operator that has been explored in earlier research [Castagna 1997; Dunfield 2014]. In addition to overloading, we can express multi-field records by merging single-field records (as already mentioned in the introduction).

Compositional programming and nested composition. Recent research on the merge operator shows that it also enables *first-class classes/traits* [Bi and Oliveira 2018] and *compositional programming* [Zhang et al. 2021]. Compositional programming supports extensible forms of datatypes and functions and offers a natural solution to hard modularity challenges, such as the Expression Problem [Wadler 1998]. At the heart of compositional programming is a mechanism, called *nested composition* [Bi et al. 2018], that composes behavior from multiple components in a merge. Nested composition differs from overloading or simple record projection, which select only one of the components in a merge. With nested composition we can write:

$$\textbf{let } mkStudent = (\lambda n : \textbf{Int}. \{age \mapsto n\}) \,\textrm{\tiny{||}}\, (\lambda n : \textbf{Int}. \{idNumber \mapsto \ldots\}) \textbf{ in } mkStudent \ 25$$

In this case we combine two functions with a merge to get an expression with the type:

$$(\textbf{Int} \rightarrow \{age : \textbf{Int}\}) \sqcap (\textbf{Int} \rightarrow \{idNumber : \textbf{Int}\})$$

Using subtyping, *mkStudent* can be given the function type $\textbf{Int} \rightarrow \{age : \textbf{Int}, idNumber : \textbf{Int}\}$; operationally the semantics of merge combines the two functions. Thus, we can use *mkStudent* to build a new record that has both an *age* and a *idNumber* field.

Merges, Ambiguity, and Subtyping. The interaction between subtyping and the merge operator is subtle. To illustrate the issue, we use an example similar to one given by Cardelli and Mitchell [1990]:

$$e_\textrm{\tiny{||}} = \textbf{let } x : \{l_2 : \textbf{Bool}\} = \{l_1 \mapsto 1, l_2 \mapsto \textbf{true}\} \textbf{ in } (\{l_1 \mapsto 2\} \,\textrm{\tiny{||}}\, x).l_1 + 3$$

In this program, x has type $\{l_2 : \textbf{Bool}\}$, despite also including a field l_1. The field l_1 is hidden due to subtyping, because $\{l_1 : \textbf{Int}, l_2 : \textbf{Bool}\} \leq \{l_2 : \textbf{Bool}\}$. The merge $\{l_1 \mapsto 2\} \,\textrm{\tiny{||}}\, x$ appears to be safe, statically, because the type of x does not contain l_1. However, what should happen when we project l_1? If the original field l_1 is preserved in x then, when we later lookup l_1, there will be two l_1 fields. If we use a biased lookup, which returns either the first value from the left or the first value from the right in a merge, then the semantics of programs may lead to surprising behaviour. For instance, in the program above, if a right-biased lookup is used, then the program would return 4. However, a programmer may have expected 5 as a result, because the type of x appears to promise that no field l_1 is present. Moreover, if the types of the two l_1 fields are distinct, this program could lead to a runtime type-error (when the field of the wrong type is projected), unless special care in taken to prevent such a situation. In essence, we would like that information hidden via subtyping has no effect in later uses of values with hidden information. For this reason Cardelli and Mitchell argued that biased lookups should *not* be used. More detailed discussions about such issues can be found in work by Huang et al. [2021].

Disjoint Intersection Types. To address the ambiguity problems, as well as the problems arising from the interactions between merges and subtyping, Oliveira et al. [2016] proposed to restrict merges so that only mergeable types are accepted. Disjointness rejects ambiguous programs such as $\textbf{true} \,\textrm{\tiny{||}}\, \textbf{false}$, because the types of the two values being merged are not disjoint. Moreover, disjointness ensures that the merge operator is symmetric (or unbiased), guaranteeing both the associativity and commutativity of the operator.

Originally, the semantics of the merge operator with disjoint intersection types was defined by elaboration, following the approach promoted by Dunfield [2014]. More recently, Huang et al. [2021] proposed a type-directed operational semantics. This approach gives a direct operational semantics to λ_i, which is a calculus with disjoint intersection types and the merge operator. In λ_i, programs can reduce without encountering ambiguities in the merges.

With a type-directed operational semantics, types are relevant at runtime, and they are used to enforce the information hiding promised by subtyping. Consider again $e_\textrm{\tiny{||}}$, the program defined

previously. In this case, we drop the field l_1 in x when the value is cast to the type $\{l_2 : \textbf{Bool}\}$. Therefore, $(\{l_1 \mapsto 2\} \mathbin{\|} x).l_1$ would become $(\{l_1 \mapsto 2\} \mathbin{\|} \{l_2 \mapsto \textbf{true}\}).l_1$ and the final result of the program would be 5. In other words, this solution to the problem of the interaction between merges and subtyping ensures that components of a merge that are hidden by subtyping are no longer accessible from the value after upcasting.

While the existing approaches to disjointness can deal with programs that have merges of records or that use nested composition, they have restricted support for overloading. For instance, the merge used in the definition of the overloaded function f (i.e. *isDigit* $\mathbin{\|}$ *not*), would be rejected. In essence in the notion of disjointness proposed by Oliveira et al. [Oliveira et al. 2016], two functions are disjoint if their return types are disjoint. However $\textbf{Char} \rightarrow \textbf{Bool}$ and $\textbf{Bool} \rightarrow \textbf{Bool}$ have overlapping return types. The disjointness restriction contrasts with traditional approaches with overloading, where distinct input types are used to eliminate possible ambiguities for overloaded functions. In addition, none of the existing calculi with disjoint intersection types include union types, which introduce new ambiguity issues.

2.2 Challenges for Deterministic Merges with Overloading and Union Types

Union Types. In prior approaches, as exemplified by λ_i, the term $1 \mathbin{\|} \textbf{true}$ is a well-typed merge. This is because the program contexts that can eliminate an integer are distinct from boolean program contexts. For example, an integer context might be $[\cdot] + 3$, where $[\cdot]$ is a "hole" into which an integer value can be filled, but this context can never be confused with any boolean context, such as, **if** $[\cdot]$ **then** e_1 **else** e_2. Therefore, no matter how a context uses the merged value, it is unambiguous whether the 1 or **true** must be projected. For instance, we have $(1 \mathbin{\|} \textbf{true}) + 3$ evaluates to $1 + 3$.

Now consider the union type $\textbf{Int} \sqcup \textbf{Bool}$. We assume that unions are *untagged*, meaning that, at runtime, there is no extra information added to indicate whether a value of this type has the type on the left side or the right side of the union. Nevertheless, as long as integer and boolean values have distinct runtime representations, i.e., they have separate constructors, it is possible to have a construct that can tell values of one type from the other as in this example from the introduction:

$$\textbf{match } e \textbf{ with } \{(y : \textbf{Bool}) \mapsto \textbf{false}, (x : \textbf{Int}) \mapsto \textbf{true}\}$$

Given that, even without union types, pattern matching on boolean and integer values is a sensible operation, the behavior of the above match expression should be no surprise. It is, after all, equivalent to the (large) expression:

$$\textbf{match } e \textbf{ with } \{\textbf{true} \mapsto \textbf{false}, \textbf{false} \mapsto \textbf{false}, 0 \mapsto \textbf{true}, -1 \mapsto \textbf{true}, 1 \mapsto \textbf{true}, \ldots\}$$

However, now that we have introduced a single elimination form that works on the union type $\textbf{Int} \sqcup \textbf{Bool}$, merging becomes nondeterministic, as we saw in the case $e = 1 \mathbin{\|} \textbf{true}$. In other words, in λ_i, **Int** and **Bool** are not distinguishable because there is no context that can take a value that is either an **Int** or **Bool** and determine which type of value was provided. As a result, it is safe for these types to be mergeable in that setting. With unions and their elimination forms, this is no longer the case. Although permitted by λ_i, merges like $1 \mathbin{\|} \textbf{true}$ are not required for many practical applications, including ours. Thus in $F_\bowtie$ we prefer to consider **Int** and **Bool** distinguishable in exchange for sacrificing the ability to merge them.

In contrast, the type $(\textbf{Int} \rightarrow \textbf{Bool}) \sqcup (\textbf{Bool} \rightarrow \textbf{Bool})$ is quite different. Unlike integer and boolean values, functions generally do not support pattern matching; they are eliminated via application. Thus, from a type-theoretic point of view, a program like

$$\textbf{match } e \textbf{ with } \{(x : \textbf{Int} \rightarrow \textbf{Bool}) \mapsto \textbf{true}, (y : \textbf{Bool} \rightarrow \textbf{Bool}) \mapsto \textbf{false}\} \tag{1}$$

is unusual. It represents an operation on a function—*other than application*—that provides information about which inputs the function accepts. Unlike the previous example with **Int** $\sqcup$ **Bool**, there is no similar way to express this kind of typecase analysis with more primitive patterns.

Implementing such a matching construct would require that functional values are tagged with type information at runtime. Thanks to subtyping, the type tag of a value could not be directly compared with the type annotation in a pattern. Therefore, runtime execution of the subtyping algorithm would be necessary. Such an arrangement certainly has precedent in the literature [Castagna et al. 1995] and in OOP language implementations, but $F_{\bowtie}$ aims to introduce union types *without* changing the meaning of existing types. In other words, support for union types should not depend on adding new operations on the values of other types. Our approach requires neither additional type tags at runtime nor any type-tag comparison in the dispatch procedure. As we shall see, $F_{\bowtie}$ will exploit type annotations on the merge operator at runtime, but those are part of the merge construct—they are not part of the representation of values.

Overloading. Overloading faces the same issues with determinism as pattern matching on unions. When the merge operator overloads functions, we can rewrite the problematic example as:

$$((\lambda x : \textbf{Int. true}) \, \| \, (\lambda y : \textbf{Bool. false})) \, (1 \, \| \, \textbf{true})$$

The merged function can be given the type $(\textbf{Int} \to \textbf{Bool}) \sqcap (\textbf{Bool} \to \textbf{Bool})$, which, due to subtyping, is equivalent to $(\textbf{Int} \sqcup \textbf{Bool}) \to \textbf{Bool}$. Semantically, these merged functions act like the match expression that we already saw:

$$\textbf{match } e \textbf{ with } \{(x : \textbf{Int}) \mapsto \textbf{true}, (y : \textbf{Bool}) \mapsto \textbf{false}\}$$

Indeed, for this reason, $F_{\bowtie}$ unifies the syntax for pattern match expressions with the syntax for (potentially merged) lambda abstractions—they are the same thing.

2.3 Information Hiding and Pattern Matching

Overloaded Functions and Copattern Matching. A common notation to represent overloaded functions in core calculi is with λ abstractions containing a case for each overload such as:

$$\lambda\{((x : \textbf{Int}) \hookrightarrow \textbf{Bool}) \mapsto \textbf{true}, ((x : \textbf{Bool}) \hookrightarrow \textbf{Bool}) \mapsto \textbf{false}\} \qquad (2)$$

This is the normalized form of the merge we saw above. In each case, the type to the right of the $\hookrightarrow$ is the type the function returns when that case is matched.

This notation is similar to GHC's LambdaCase and OCaml's function syntax, both of which combine the introduction of a function with pattern matching on its argument. It may also be seen as a form of *copattern matching* [Abel et al. 2013; Zeilberger 2008] and is common in the literature on overloading [Castagna et al. 1995].

Typical pattern matching involves destructuring a value according to a pattern describing the shape of an introduction form. In example (2), on the other hand, the λ-value destructures its *evaluation context* according to *elimination patterns* (i.e., copatterns) that describe the shape of elimination forms. In that example, the context $[\cdot] \, 1$, matches the first elimination pattern $(x : \textbf{Int}) \hookrightarrow \textbf{Bool}$, whereas the context $[\cdot] \, \textbf{true}$ matches the second. In this way, *dispatch*—the process of deciding which overload of a function to execute—is completely subsumed by (co)pattern matching.

η-Expansion. Recall the definition of $e_{\|}$.

$$e_{\|} = \textbf{let } x : \{l_2 : \textbf{Bool}\} = \{l_1 \mapsto 1, l_2 \mapsto \textbf{true}\} \textbf{ in } (\{l_1 \mapsto 2\} \, \| \, x).l_1 \, + \, 3$$

Previously, we saw that the expected semantics should hide the l_1 field at the assignment of x. Type annotations ought to have the effect at runtime of hiding information (such as a field of a record

or an overload of a function) in a term. $F_{\bowtie}$ implements this by η-expanding annotated terms. An expression $(v : \{l_2 : \textbf{Bool}\})$ η-expands to $\{l_2 \mapsto (v.l_2 : \textbf{Bool})\}$, hiding any information, other than the contents of the field l_2, that may be present in v.

In this way, η-expansion of a value at a type builds a wrapper that limits access to the value to the operations supported by the type. This technique is quite similar to how sound gradual type systems [Siek and Taha 2006] build wrappers to catch runtime type errors caused by untyped code passing ill-typed values to typed contexts.

Formally, in $F_{\bowtie}$, a value v inside a wrapper for type A is written $\lambda\{(* \hookrightarrow A) \mapsto v\}$, so we would write the example above as $\lambda\{(* \hookrightarrow \{l_2 : \textbf{Bool}\}) \mapsto v\}$. Here, the elimination pattern $* \hookrightarrow A$ matches any context valid for type A. When this occurs, v is placed inside this context. In other words, the value $\lambda\{(* \hookrightarrow A) \mapsto v\}$ behaves exactly as v but only in contexts valid for type A. The $*$ pattern is a sort of dual to variable patterns.

Pattern Expansion is η-Expansion. We have already seen that type annotations on pattern variables have significance at runtime. Their meaning is derived from η principals for positive types. Consider an $F_{\bowtie}$ expression like **match** e **with** $\{(x : \textbf{Bool}) \mapsto e'\}$. Dually to η expansion for a value of record type, which re-builds an output and projects each field individually, η expansion for a positive type like **Bool** re-builds an input by matching against all possible values of a type and specializing each branch. Given the term e' above has a free variable $x : \textbf{Bool}$, the general form of η expansion for booleans would be:

$$e' \rightarrow_\eta \textbf{match } x \textbf{ with } \{\textbf{true} \mapsto e'[\textbf{true}/x], \textbf{false} \mapsto e'[\textbf{false}/x]\}$$

From this, one might informally expect an equivalence:

$$\textbf{match } e \textbf{ with } \{(x : \textbf{Bool}) \mapsto e'\} \equiv \textbf{match } e \textbf{ with } \{\textbf{true} \mapsto e'[\textbf{true}/x], \textbf{false} \mapsto e'[\textbf{false}/x]\}$$

In $F_{\bowtie}$, this expansion is not just an equivalence: it is the definition of the operational semantics of pattern matching. This is possible because similar expansions exist for all positive types. During elimination pattern matching, a dual kind of expansion occurs for $*$-patterns. The wrapper $\lambda\{(* \hookrightarrow \{l_2 : \textbf{Bool}\}) \mapsto v\}$ is equivalent to $\{l_2 \mapsto (v.l_2 : \textbf{Bool})\}$. In other words, the pattern $* \hookrightarrow A^-$ η-expands according to the negative type A^-. In summary, in $F_{\bowtie}$ we can characterize the type-directed component of the operational semantics as runtime η-expansion.

2.4 Compositional Programming in $F_{\bowtie}$

To see how $F_{\bowtie}$ supports the kind of (nested) compositional programming offered by the merge operator, and to introduce the notation used by the calculus, we next consider how to build a small extensible interpreter.

In $F_{\bowtie}$, we can represent a language of integer literals and addition expressions by defining the type **IntArithExpr** A as a record type as:

$$\textbf{IntArithExpr } A = \{constant : \textbf{Int} \rightarrow A, add : A \rightarrow A \rightarrow A\}$$

This type represents a simplified form of a *compositional interface* [Zhang et al. 2021], and is closely related to the kind of interfaces used in techniques such as *finally tagless* [Carette et al. 2009] or *object algebras* [Oliveira and Cook 2012]. Indeed this interface is essentially the fold (F-)Algebra [Bird and de Moor 1996] for a simple datatype of arithmetic expressions. An $F_{\bowtie}$ term of type **IntArithExpr** A describes how to interpret an expression in our object language as a value of type A. We represent one such object-language expression as follows:

> $three : \forall \alpha.\textbf{IntArithExpr } \alpha \rightarrow \alpha$
> $three = \Lambda\alpha.\lambda(x : \textbf{IntArithExpr } \alpha) \hookrightarrow \alpha. \; x.add \; (x.constant \; 1) \; (x.constant \; 2)$

The notation $\Lambda\alpha.e$ binds a type variable α in the body e while $\lambda(x : B) \hookrightarrow C.\, e$ binds the term variable x of type B in the body e, which is ascribed the type C. Values such as *three* are defined by implementing their folds.

A natural way of interpreting object-language expressions as integers is by evaluating them. The interpretation *evalInt* describes how to do this, by defining a meaning for each field:

$$evalInt \;:\; \textbf{IntArithExpr Int}$$
$$evalInt = \{constant \mapsto \lambda x : \textbf{Int} \hookrightarrow \textbf{Int}.\; x,\, add \mapsto (+)\}$$

Now the program *three* [**Int**] *evalInt*, of type **Int**, evaluates to 3.

Suppose that we wish to extend this language so that a constant may be either an integer or a floating point number. For simplicity, we will not presume any subtyping relationship between **Int** and **Float**, but instead rely on an explicit cast *toFloat* : **Int** → **Float**. What we wish is to obtain a combined language where both floating point numbers and integers can be used. Moreover, the language should automatically convert between integers and floating point numbers when necessary. Orchard and Schrijvers [2010] tackle a similar problem in the setting of a typed object language. They illustrate that the problem is tricky to solve in Haskell using a finally tagless embedding. To this end, Orchard and Schrijvers proposed to extend Haskell with *constraint synonyms*, which later helped motivate the addition of the ConstraintKinds GHC extension. In contrast, in $F_{\bowtie}$, merge-based overloading with intersection and union types provides what we need.

We define the type **IntFloatArithExpr** A as an extension of **IntArithExpr** A using intersection. Observe that the intersection distributes over the record, augmenting the type of the *constant* field.

$$\textbf{IntFloatArithExpr } A = \textbf{IntArithExpr } A \;\sqcap\; \{constant : \textbf{Float} \to A\}$$
$$\equiv \{constant : (\textbf{Int} \to A) \sqcap (\textbf{Float} \to A),\, add : A \to A \to A\}$$
$$\equiv \{constant : \textbf{Int} \sqcup \textbf{Float} \to A,\, add : A \to A \to A\}$$

In other words, subtyping gives us the equivalence $(\textbf{Int} \to A) \sqcap (\textbf{Float} \to A) \equiv (\textbf{Int} \sqcup \textbf{Float}) \to A$. Such equivalences are key to compositional programming [Bi et al. 2019].

To evaluate expressions in this language, we need to define an interpretation *evalIntFloat* : **IntFloatArithExpr** (**Int** ⊔ **Float**). A first step is to define the type expressions as well as the evaluation function for the sub-language containing floating point but not integer literals. (We use +. for the floating-point addition primitive.)

$$\textbf{FloatArithExpr } A = \{constant : \textbf{Float} \to A,\, add : A \to A \to A\}$$

$$evalFloat \;:\; \textbf{FloatArithExpr Float}$$
$$evalFloat = \{constant \mapsto \lambda x : \textbf{Float} \hookrightarrow \textbf{Float}.\; x,\, add \mapsto (+.)\}$$

It is now possible to merge our two evaluators.

$$partialEvalIntFloat \;:\; (\textbf{IntArithExpr Int}) \sqcap (\textbf{FloatArithExpr Float})$$
$$partialEvalIntFloat = evalInt \;,\!,\; evalFloat$$

But, we are not finished: we would like an expression of type **IntFloatArithExpr** (**Int** ⊔ **Float**), but only have one of type (**IntArithExpr Int**)⊓(**FloatArithExpr Float**). This is not enough because the former type requires the *add* operation to support addition of integers with floating point numbers, which the latter does not. As a result, trying to use *partialEvalIntFloat* to evaluate *four* (defined below) would be ill typed.

$$four = \Lambda\alpha.\lambda(x : \textbf{IntFloatArithExpr } \alpha) \hookrightarrow \alpha.$$
$$x.add\; (x.constant\; 2.0)\; (x.constant\; 2)$$

expressions	e	::=	$x \mid c\,e \mid e_1\,e_2 \mid e\,[A]$	*types*	A, B, C	::=	$\alpha \mid c\,A$
			$\mid (e : A)$				$\mid A \to B \mid \forall \alpha.B$
			$\mid (e_1 : A_1) \parallel (e_2 : A_2)$				$\mid \bot \mid A_1 \sqcup A_2$
			$\mid \lambda\{\hat{p}_1 \mapsto e_1, ..., \hat{p}_n \mapsto e_n\}$				$\mid \top \mid A_1 \sqcap A_2$
				neg. types	A^-, B^-	::=	$A \to B \mid \forall \alpha.B$
values	v	::=	$x \mid c\,v$				$\mid A_1^- \sqcup A_2^-$
			$\mid \lambda\{\hat{p}_1 \mapsto e_1, ..., \hat{p}_n \mapsto e_n\}$				$\mid \top \mid A_1^- \sqcap A_2^-$
value patterns	p	::=	$x : A \mid c\,p \mid (p_1 \mid p_2)$				
			$\mid p_1 \& p_2$	*elim. types*	$\hat{A}, \hat{B}$	::=	$A \mid [A]$
elim. frames	F	::=	$[\cdot]v \mid [\cdot][A]$	*environments*	Γ, Δ	::=	$\cdot \mid \Gamma, \alpha \mid \Gamma, x : A$
elim. patterns	$\hat{p}$	::=	$* \hookrightarrow A^- \mid p \hookrightarrow B \mid \alpha \hookrightarrow B$				

Fig. 1. $F_{\bowtie}$ Syntax

In order to complete the evaluator, we need an addition extension that handles the missing cases. Its type is given below.

$$\mathbf{ArithExt} = \{add : (\mathbf{Int} \to \mathbf{Float} \to \mathbf{Float}) \sqcap (\mathbf{Float} \to \mathbf{Int} \to \mathbf{Float})\}$$

We implement the extension by coercing integers to floats and using floating point addition, as shown in the code below. Note that the *add* field of the record contains a "function" with two bodies, distinguished by the input type of the argument x.

$$
\begin{aligned}
evalExt \;&:\; \mathbf{ArithExt} \\
evalExt = \{add \mapsto \lambda\{\; &(x : \mathbf{Int}) \mapsto \lambda(y : \mathbf{Float}) \hookrightarrow \mathbf{Float}.\; toFloat\; x +.\; y, \\
&(x : \mathbf{Float}) \mapsto \lambda(y : \mathbf{Int}) \hookrightarrow \mathbf{Float}.\; x +.\; toFloat\; y \;\} \;\}
\end{aligned}
$$

Applying distributivity, we can show that

$$(\mathbf{IntArithExpr\ Int}) \sqcap (\mathbf{FloatArithExpr\ Float}) \sqcap \mathbf{ArithExt} \le \mathbf{IntFloatArithExpr}\ (\mathbf{Int} \sqcup \mathbf{Float})$$

This tells us that merging *partialEvalIntFloat* with *evalExt* yields an expression of the desired type.

$$
\begin{aligned}
evalIntFloat \;&:\; \mathbf{IntFloatArithExpr}\ (\mathbf{Int} \sqcup \mathbf{Float}) \\
evalIntFloat \;&=\; partialEvalIntFloat \parallel evalExt
\end{aligned}
$$

Now, evaluating *four* [$\mathbf{Int} \sqcup \mathbf{Float}$] *evalIntFloat* results in the floating-point value 4.0 as expected.

3　SYSTEM $F_{\bowtie}$

We will now make the intuitions from earlier precise. This section presents the syntax and operational semantics of $F_{\bowtie}$ in full.

3.1　Syntax

Expressions. Figure 1 presents the syntax of $F_{\bowtie}$. Expressions include standard syntax such as variables, application of functions to arguments, and application of terms to types. An underscore is used in place of a variable name when the variable is never referenced. The $c\,e$ form, tags the expression e with constructor c. The constructor c comes from some predetermined set of symbols; the metavariables l and k also range over this set, particularly when the constructor represents a record label. The notation $(e : A) \parallel (e' : B)$ denotes a merge of two type-annotated expressions.

　　We denote a subset of expressions as values: variables, constructors with values as arguments, and λ forms. We also describe *elimination frames* which are contexts in the shape of an elimination form with a hole in head position. Elimination frames play a role dual to that of values in pattern matching.

$$A + B = \text{left } A \sqcup \text{right } B \qquad \textbf{Bool} = \top + \top \qquad \{l : A\} = l \top \to A$$
$$\{l_1 : A_1, ..., l_n : A_n\} = \{l_1 : A_1\} \sqcap ... \sqcap \{l_n : A_n\}$$

$$\textbf{true} = \text{left } \lambda\{\} \qquad \textbf{false} = \text{right } \lambda\{\} \qquad e.l = e\,(l\,\lambda\{\})$$
$$\lambda x : A.\ e = \lambda\{x : A \mapsto e\} \qquad \Lambda\alpha.e = \lambda\{\alpha \mapsto e\}$$
$$\{l_1 \mapsto e_1, ..., l_n \mapsto e_n\} = \lambda\{(\_ : l_1 \top) \mapsto e_1, ..., (\_ : l_n \top) \mapsto e_n\}$$
$$\textbf{match } e \textbf{ with } \{p_1 \mapsto e_1, ..., p_n \mapsto e_n\} = \lambda\{p_1 \mapsto e_1, ..., p_n \mapsto e_n\}\ e$$

Fig. 2. Useful Abbreviations

Types. The types of $F_{\bowtie}$ include type variables α, constructors applied to an argument $c\,A$, functions $A \to B$, and universal quantifiers $\forall\,\alpha.B$. Additionally, the $\bot$ type is uninhabited while $\top$ is inhabited by the trivial value $\lambda\{\}$. Finally, $F_{\bowtie}$ types include unions, $A \sqcup B$, and intersections, $A \sqcap B$. We sometimes write n-ary intersections and unions such as $A_1 \sqcap ... \sqcap A_n$. This notation does not preclude the possibility that $n = 0$, in which case the type is $\top$ (or $\bot$ in the case of 0-ary unions). We will in §4.2 introduce some restrictions on which unions and intersections are considered to be well-formed.

Negative types are those that are introduced by λ and eliminated with strong (non-matching) elimination forms. These include functions, quantifiers, $\top$, and unions and intersections of negative types. Elimination types are associated with elimination forms. An elimination type $\hat{A} = A$ corresponds to an application $[\cdot]\,e$ in which the argument e has type A; the elimination type $\hat{A} = [A]$ corresponds to a polymorphic instantiation $[\cdot][A]$.

The typing algorithm makes use of unordered environments Γ which both contain type variables and map term variables to their types. The type associated with a term variable in an environment Γ may reference any free type variable in Γ.

Encodings. Various standard syntax can be encoded in terms of $F_{\bowtie}$'s constructs. For example, one might in theory encode **Int**, the type of 32-bit integers, as a union of 2^{32} distinct constructors. Figure 2 contains other useful abbreviations. This figure omits (as we often do) return type annotations $\hookrightarrow A$ where they are to be inferred from context. Sum types are represented by unions of the presumed-distinct left and right constructors. Record types are also encodable as functions from field labels to the corresponding value. The type of a label is represented as $l \top$ (the constructor l with a dummy argument of type $\top$). A single-field record is a function accepting only the label $l\,\lambda\{\}$ as input.

Patterns. Pattern matching is a distinguishing feature of $F_{\bowtie}$. There are two types of patterns. The metavariable p ranges over *value patterns*, which are essentially the patterns of Haskell or ML. In a match construct **match** v **with** $\{p_1 \mapsto e_1, ..., p_n \mapsto e_n\}$, the discriminee v will be deconstructed by the patterns of each $p_i \mapsto e_i$ clause.

As we have seen, the λ form contains a number of overloaded implementations of a computation, with the correct one chosen dynamically based on how the value is used by its context. Each overload is represented as a clause $\hat{p} \mapsto e$ where $\hat{p}$ is an elimination pattern (or copattern [Abel et al. 2013; Zeilberger 2008]). We can think of λ as using these patterns to build a computation by deconstructing its immediate context—specifically, the elimination frame F. For example, consider the following value:

$$\lambda\{((x : A_1) \hookrightarrow B_1) \mapsto e_1, ((x : A_2) \hookrightarrow B_2) \mapsto e_2, (\alpha \hookrightarrow B_3) \mapsto e_3\}$$

In a frame of shape $[\cdot]\,v'$, the e_1 or e_2 overloads are selected when v' has type A_1 or A_2 respectively. (In fact, both may be be selected and merged together if A_1 and A_2 are not distinguishable.) In a

$$\boxed{e \longmapsto e'} \qquad\qquad\qquad\qquad \textit{(small-step reduction)}$$

$$
\begin{aligned}
M &::= \{\hat{p}_1 \mapsto e_1, \dots, \hat{p}_n \mapsto e_n\} \\
E &::= [\cdot] \mid c\,E \mid E\,e \mid v\,E \mid E[A] \\
&\quad\mid (E : A) \mid (E : A) \,\|\, (e : B) \\
&\quad\mid (v : A) \,\|\, (E : B)
\end{aligned}
$$

$$\frac{e \longmapsto e'}{E[e] \longmapsto E[e']} \qquad \frac{}{F[\lambda M] \longmapsto \mathrm{disp}(M, F)} \qquad \frac{}{(v : A) \longmapsto v}$$

$$\frac{}{(v_1 : A_1^-) \,\|\, (v_2 : A_2^-) \longmapsto \lambda\{* \hookrightarrow A_1^- \mapsto v_1, * \hookrightarrow A_2^- \mapsto v_2\}}$$

$$\mathrm{disp}(M, F) = \underset{\substack{\hat{p} \mapsto e \in M \\ F/\hat{p} \Rightarrow \sigma \hookrightarrow B}}{\Big\|}\ (\sigma(e) : B)$$

Fig. 3. $F_{\bowtie}$ Operational Semantics and Auxiliary Definitions

context of shape $[\cdot][A]$, the third clause is matched instead. Each B_i in the value is an output type annotation that describes the type of the corresponding expression e_i. Note that α is bound in B_3.

A clause $(* \hookrightarrow A^-) \mapsto v$ matches any elimination F that is valid for a value of type A^-. The elimination frame F, then becomes the context for v. In this way, the elimination pattern $*$ is a sort of "pattern variable" that places the current context around the expression in which $*$ is "bound". The point is purely to enforce that v has precisely the type A^-. This may seem unusual, but the semantics are straightforward. A term $\lambda\{(* \hookrightarrow \{l_1 : A_1\}) \mapsto \{l_1 \mapsto e_1, l_2 \mapsto e_2\}\}$ should be thought of as a wrapper around $\{l_1 \mapsto e_1, l_2 \mapsto e_2\}$ that enforces the type $\{l_1 : A_1\}$ by hiding the field l_2.

As we have seen, this hiding is achieved using η-expansion. In other words, we are taking advantage of the fact that $\lambda\{(* \hookrightarrow \{l_1 : A\}) \mapsto v\}$ is equivalent to $\lambda\{(x : l_1 \top) \hookrightarrow A \mapsto v.l_1\}$. This explains why A^- must be a negative type: only negative types have more primitive elimination patterns to introduce them with. The syntax of value patterns also includes or- and and-patterns written $p_1 | p_2$ and $p_1 \& p_2$. These provide η-principles for union and intersection types. Variables in patterns are annotated with the type of value they are expected to match. They are bound in the right hand side of a clause.

3.2 Operational Semantics

A small-step, call-by-value operational semantics is given in Figure 3. We use M to range over λ-bodies and E to range over evaluation contexts. Evaluation proceeds under all evaluation contexts as is standard. The metavariable σ ranges over substitutions: mappings from term and type variables to values and types. A substitution may also contain at most one special mapping from $*$ to an elimination frame. The notation $\sigma(e)$ applies the substitution: every variable in the domain of σ is replaced with the corresponding type or value in e. The result is then placed in the frame F if such a frame exists in σ.

The usual β rules for type and term application are subsumed into a single rule which evaluates terms of the form $F[\lambda M]$ (strong elimination forms applied to a value) using the dispatch meta-function, disp. This procedure looks at every case $\hat{p} \mapsto e$ in M and checks whether F matches $\hat{p}$. If so, a substitution σ and output type B are obtained. The result of disp is then a merge of the expression e of every matching case with an appropriate substitution applied for pattern variables. The notation $F/\hat{p} \Rightarrow \sigma \hookrightarrow B$ used here indicates that elimination frame F matches elimination pattern $\hat{p}$; we will shortly present this process in more detail. A merge of two values evaluates to a λ-expression with two corresponding cases. Note that the type that each of the values is annotated with is remembered; the λ acts as a wrapper to hide access to any part of either value not described by its type annotation. Type annotations on expressions other than merges can be safely dropped during evaluation, since if the expression gets merged it will be annotated again in the merge.

$\boxed{v/p \Rightarrow \sigma}$ *(value pattern matching)*

$$\frac{v/p \Rightarrow \sigma}{c\,v/c\,p \Rightarrow \sigma} \qquad \frac{v/p_1 \Rightarrow \sigma}{v/(p_1|p_2) \Rightarrow \sigma} \qquad \frac{v/p_2 \Rightarrow \sigma}{v/(p_1|p_2) \Rightarrow \sigma} \qquad \frac{v/p_1 \Rightarrow \sigma_1 \qquad v/p_2 \Rightarrow \sigma_2 \qquad \sigma = \sigma_1 \sqcap \sigma_2}{v/(p_1\&p_2) \Rightarrow \sigma}$$

$$\frac{}{v/(x : A^-) \Rightarrow [x \mapsto v]} \qquad \frac{v/((x : A_1)\&(x : A_2)) \Rightarrow \sigma}{v/(x : A_1 \sqcap A_2) \Rightarrow \sigma} \qquad \frac{v/c\,(x : A) \Rightarrow [x \mapsto v']}{v/(x : c\,A) \Rightarrow [x \mapsto c\,v']}$$

$$\frac{v/((x : A_1)|(x : A_2)) \Rightarrow \sigma}{v/(x : A_1 \sqcup A_2) \Rightarrow \sigma}$$

$\boxed{F/\hat{p} \Rightarrow \sigma \hookrightarrow B}$ *(elimination pattern matching)*

$$\frac{v/p \Rightarrow \sigma}{[\cdot]v/(p \hookrightarrow B) \Rightarrow \sigma \hookrightarrow B} \qquad \frac{}{[\cdot][A]/(\alpha \hookrightarrow B) \Rightarrow [\alpha \mapsto A] \hookrightarrow B[A/\alpha]}$$

$$\frac{F/((x : A) \hookrightarrow B) \Rightarrow [x \mapsto v] \hookrightarrow B'}{F/(* \hookrightarrow (A \to B)) \Rightarrow [* \mapsto [\cdot]v] \hookrightarrow B'} \qquad \frac{F/(\alpha \hookrightarrow B) \Rightarrow [\alpha \mapsto A'] \hookrightarrow B'}{F/(* \hookrightarrow \forall \alpha.B) \Rightarrow [* \mapsto [\cdot][A']] \hookrightarrow B'}$$

$$\frac{F/(* \hookrightarrow A_1^-) \Rightarrow \sigma_1 \hookrightarrow B_1 \qquad}{F/(* \hookrightarrow A_2^-) \Rightarrow \sigma_2 \hookrightarrow B_2 \qquad \sigma = \sigma_1 \sqcap \sigma_2} \qquad \frac{F/(* \hookrightarrow A_1^-) \Rightarrow \sigma_1 \hookrightarrow B_1 \qquad}{F/(* \hookrightarrow A_2^-) \Rightarrow \sigma_2 \hookrightarrow B_2 \qquad \sigma = \sigma_1 \sqcup \sigma_2}$$
$$\overline{F/(* \hookrightarrow A_1^- \sqcup A_2^-) \Rightarrow \sigma \hookrightarrow B_1 \sqcup B_2} \qquad \overline{F/(* \hookrightarrow A_1^- \sqcap A_2^-) \Rightarrow \sigma \hookrightarrow B_1 \sqcap B_2}$$

$$\frac{F/(* \hookrightarrow A_1^-) \Rightarrow \sigma \hookrightarrow B \qquad F/(* \hookrightarrow A_2^-) \Rightarrow\!\!\!\!/}{F/(* \hookrightarrow A_1^- \sqcap A_2^-) \Rightarrow \sigma \hookrightarrow B} \qquad \frac{F/(* \hookrightarrow A_1^-) \Rightarrow\!\!\!\!/ \qquad F/(* \hookrightarrow A_2^-) \Rightarrow \sigma \hookrightarrow B}{F/(* \hookrightarrow A_1^- \sqcap A_2^-) \Rightarrow \sigma \hookrightarrow B}$$

$\boxed{\sigma = \sigma_1 \sqcap \sigma_2}$ $\qquad\qquad$ $\boxed{\sigma = \sigma_1 \sqcup \sigma_2}$

$(\sigma_1 \sqcap \sigma_2)(\alpha)$	$=$	A_1	where $\sigma_1(\alpha) = A_1$, α not in σ_2		$(\sigma_1 \sqcup \sigma_2)(\alpha)$	$=$	A	
$(\sigma_1 \sqcap \sigma_2)(\alpha)$	$=$	A_2	where $\sigma_2(\alpha) = A_2$, α not in σ_1					where $\sigma_1(\alpha) = \sigma_2(\alpha) = A$
$(\sigma_1 \sqcap \sigma_2)(\alpha)$	$=$	A	where $\sigma_1(\alpha) = \sigma_2(\alpha) = A$		$(\sigma_1 \sqcup \sigma_2)(x)$	$=$	v	
								where $\sigma_1(x) = \sigma_2(x) = v$
$(\sigma_1 \sqcap \sigma_2)(x)$	$=$	v_1	where $\sigma_1(x) = v_1$, x not in σ_2		$(\sigma_1 \sqcup \sigma_2)(*)$	$=$	F	
$(\sigma_1 \sqcap \sigma_2)(x)$	$=$	v_2	where $\sigma_2(x) = v_2$, x not in σ_1					where $\sigma_1(*) = \sigma_2(*) = F$
$(\sigma_1 \sqcap \sigma_2)(x)$	$=$	v	where $\sigma_1(x) = \sigma_2(x) = v$					
$(\sigma_1 \sqcap \sigma_2)(*)$	$=$	F_1	where $\sigma_1(*) = F_1$, $*$ not in σ_2					
$(\sigma_1 \sqcap \sigma_2)(*)$	$=$	F_2	where $\sigma_2(*) = F_2$, $*$ not in σ_1					
$(\sigma_1 \sqcap \sigma_2)(*)$	$=$	F	where $\sigma_1(*) = \sigma_2(*) = F$					

Fig. 4. Pattern Matching

Pattern Matching. Figure 4 gives the pattern matching algorithm. We write $v/p \Rightarrow \sigma$ to mean the value v matches the value pattern p, where σ describes the corresponding bindings for pattern variables. During the matching process, constructors are compared to a pattern structurally. Or-patterns match a value when either one of the two component patterns match. The type system will ensure that the two patterns are mutually exclusive, so there is no need for a rule that handles both patterns matching. An and-pattern matches a value when both of its subpatterns match. The results of the two matches are combined with the partial $\sigma_1 \sqcap \sigma_2$ operation. On substitutions with disjoint domains, this operation is concatenation. However, when the domains overlap it is defined only when any variables present in both results are assigned the same value.

Variable patterns annotated with negative types trivially match any value because values of negative type have no discerning tags to take advantage of in the pattern matching process. On the other hand, when the type is positive we can perform η-expansion. The last three cases of the value pattern matching definition define matching of a variable pattern in terms of patterns for the corresponding type. In this way, these variable patterns are essentially just shortcuts for more primitive patterns.

A value may match a pattern variable annotated with a negative intersection or union type either via the negative variable case or via the intersection/union case. The following lemma makes clear that we get the same substitution as output either way.

LEMMA 3.1. *If $v/(x : A) \Rightarrow \sigma$ then $\sigma = [x \mapsto v]$.*

PROOF. Routine induction on A.   □

Thus, this issue is not a source of nondeterminism. However, the treatment of or-patterns does introduce nondeterminism. A value v that matches the patterns p_1 and p_2 in two different ways can also match $p_1|p_2$ in both ways. The type system will rule out such nondeterministic examples by ensuring only one side of an or-pattern ever matches. In other words, well-typed or-patterns are exclusive and pattern matching is deterministic over well-typed patterns.

Dually, $F/\hat{p} \Rightarrow \sigma \hookrightarrow B$ means that the elimination frame F matches the pattern $\hat{p}$, where the substitution σ contains bindings for pattern variables and may additionally contain an elimination frame. The negation of this definition, $F/\hat{p} \nRightarrow$ means that F does not match $\hat{p}$. The type B is the expected return type of the case of the λ-expression in which $\hat{p}$ appears.

The $*$ elimination patterns behave in a dual way to variable patterns. To understand them, it is helpful to see an example. Consider the term $e = \lambda\{(* \hookrightarrow (A \rightarrow B)) \mapsto v\}$ v'. Here, $F = [\cdot]v'$ and $\hat{p} = * \hookrightarrow A \rightarrow B$. The matching procedure proceeds by recursively matching F against the pattern $(x : A) \hookrightarrow B$. In the end, the result is $[* \mapsto [\cdot]v']$ with a return type of B. Thus, $e \longmapsto (v\ v' : B)$.

As with variable patterns for positive types, $*$ patterns for function and polymorphic types are defined in terms of more primitive patterns. To keep $F_{\bowtie}$'s typechecking deterministic, however, the language includes neither and-elimination patterns nor or-elimination patterns. Thus, $*$ patterns for union and intersection types must be treated specially. This in effect offloads the nondeterminism to the subtyping algorithm. The dual role of elimination patterns when compared with value patterns can be counterintuitive when it comes to union and intersection types. For example, a value matches a value pattern $x : A_1 \sqcup A_2$ when it matches one of either $x : A_1$ or $x : A_2$ but an elimination frame matches an elimination pattern $* \hookrightarrow A_1 \sqcup A_2$ when it matches *both* $* \hookrightarrow A_1$ and $* \hookrightarrow A_2$. Similarly, $\sigma_1 \sqcap \sigma_2$ is used to combine substitutions during elimination matching for union types while the dual $\sigma_1 \sqcup \sigma_2$ operation is used for intersection types.

4 TYPE SYSTEM

We now present the type system of $F_{\bowtie}$. It aims to guarantee both a conventional type soundness property as well as determinism of reduction.

4.1 Subtyping

Types form a bounded distributive lattice under the subtyping relation (written $A \leq B$) defined in Figure 5. The standard subtyping rules for function, union, and intersection types are included in this definition or derivable from the rules that are. For example, the rule

$$\frac{B \leq A \qquad A' \leq B'}{A \rightarrow A' \leq B \rightarrow B'}$$

$A \leq B$ *(declarative subtyping)*

$$\frac{}{A \leq A} \qquad \frac{A_1 \leq A_2 \qquad A_2 \leq A_3}{A_1 \leq A_3} \qquad \frac{A \leq B}{c\,A \leq c\,B} \qquad \frac{A \leq B}{C \to A \leq C \to B} \qquad \frac{A \leq B}{\forall\,\alpha.A \leq \forall\,\alpha.B}$$

$$\frac{A_2 \leq A_1}{A_1 \to B \leq A_2 \to B} \qquad \frac{A_1 \leq B}{A_1 \sqcap A_2 \leq B} \qquad \frac{A_2 \leq B}{A_1 \sqcap A_2 \leq B} \qquad \frac{A \leq B_1 \qquad A \leq B_2}{A \leq B_1 \sqcap B_2}$$

$$\frac{}{(C \to A) \sqcap (C \to B) \leq (C \to A \sqcap B)} \qquad \frac{}{(c\,A) \sqcap (c\,B) \leq c\,(A \sqcap B)}$$

$$\frac{}{(\forall\,\alpha.A) \sqcap (\forall\,\alpha.B) \leq (\forall\,\alpha.A \sqcap B)} \qquad \frac{}{\forall\,\alpha.(A \sqcup B) \leq (\forall\,\alpha.A) \sqcup (\forall\,\alpha.B)} \qquad \frac{}{c\,(A \sqcup B) \leq (c\,A) \sqcup (c\,B)}$$

$$\frac{}{(A_1 \to B) \sqcap (A_2 \to B) \leq (A_1 \sqcup A_2) \to B} \qquad \frac{A \leq B_1}{A \leq B_1 \sqcup B_2} \qquad \frac{A \leq B_2}{A \leq B_1 \sqcup B_2} \qquad \frac{}{\perp \leq B} \qquad \frac{}{A \leq \top}$$

$$\frac{}{(A \sqcup B) \sqcap C \leq (A \sqcap C) \sqcup (B \sqcap C)} \qquad \frac{A_1 \leq B \qquad A_2 \leq B}{A_1 \sqcup A_2 \leq B}$$

Fig. 5. Declarative Subtyping

$A \bowtie B$ *(mergeability)* $\qquad A \bowtie_{ax} B$ *(merg. axioms)*

$$\frac{A \diamond B}{A \to A' \bowtie B \to B'} \qquad \frac{B \bowtie B'}{A \to B \bowtie A \to B'} \qquad \frac{A \bowtie B \qquad A' \bowtie B}{A \sqcap A' \bowtie B} \qquad \frac{}{\top \bowtie_{ax} B}$$

$$\frac{A \bowtie B}{\forall\,\alpha.A \bowtie \forall\,\alpha.B} \qquad \frac{A \bowtie B \qquad A' \bowtie B}{A \sqcup A' \bowtie B} \qquad \frac{A \bowtie_{ax} B}{A \bowtie B} \qquad \frac{B \bowtie A}{A \bowtie B} \qquad \frac{}{A \to A' \bowtie_{ax} \forall\,\alpha.B}$$

$A \diamond B$ *(distinguishability)* $\qquad A \diamond_{ax} B$ *(dist. axioms)*

$$\frac{A \diamond B}{c\,A \diamond c\,B} \qquad \frac{A \diamond B \qquad A' \diamond B}{A \sqcup A' \diamond B} \qquad \frac{A \diamond_{ax} B}{A \diamond B} \qquad \frac{B \diamond A}{A \diamond B} \qquad \frac{}{\perp \diamond_{ax} B}$$

$$\frac{A \diamond B' \qquad B \leq B'}{A \diamond B} \qquad \frac{c_1 \neq c_2}{c_1\,A \diamond_{ax} c_2\,B}$$

Fig. 6. Disjointness Relations

is derivable from transitivity and the co- and contravariance rules for the function type constructor. Subtyping is defined to be reflexive and transitive. Additionally, unions and intersections distribute over most other types. The following type equivalences hold:

$$\begin{aligned} c\,(A \sqcup B) &\equiv (c\,A) \sqcup (c\,B) & \forall\,\alpha.(A \sqcap B) &\equiv (\forall\,\alpha.A) \sqcap (\forall\,\alpha.B) \\ (C \to A) \sqcap (C \to B) &\equiv C \to (A \sqcap B) & (A \to C) \sqcap (B \to C) &\equiv (A \sqcup B) \to C \end{aligned}$$

The notation $A \equiv B$ denotes subtyping in both directions.

Due to the presence of overloading, we do not in general have $(A \to A') \sqcap (B \to B') \leq (A \sqcup B) \to (A' \sqcap B')$. Consider an overloaded function which produces an integer when given an integer and produces a boolean when given a boolean. It has type $(\textbf{Int} \to \textbf{Int}) \sqcap (\textbf{Bool} \to \textbf{Bool})$. It cannot be cast to the type $(\textbf{Int} \sqcup \textbf{Bool}) \to (\textbf{Int} \sqcap \textbf{Bool})$ because given only one of an integer or boolean, it is able to produce a value of only one of those types as a result—not both.

$$\boxed{\Gamma \vdash A} \hspace{6cm} \textit{(type well-formedness)}$$

$$\frac{\alpha \in \Gamma}{\Gamma \vdash \alpha} \qquad \frac{\Gamma \vdash A}{\Gamma \vdash c\,A} \qquad \frac{\Gamma \vdash A \quad \Gamma \vdash B}{\Gamma \vdash A \to B} \qquad \frac{\Gamma, \alpha \vdash B}{\Gamma \vdash \forall \alpha.B} \qquad \frac{}{\Gamma \vdash \top} \qquad \frac{}{\Gamma \vdash \bot}$$

$$\frac{\Gamma \vdash A_1 \quad \Gamma \vdash A_2 \quad A_1 \bowtie A_2}{\Gamma \vdash A_1 \sqcap A_2} \qquad \frac{\Gamma \vdash A_1 \quad \Gamma \vdash A_2 \quad A_1 \,\text{\footnotesize\textbullet}\, A_2}{\Gamma \vdash A_1 \sqcup A_2}$$

Fig. 7. Well-Formed Types

4.2 Disjoint and Well-Formed Types

The mergeability and distinguishability relations are defined inductively in Figure 6.

Mergeability. The idea of mergeability is to relate two types A and B when there is no potentially ambiguous operation shared between them. Consider the following derivable rules.

$$\frac{l \neq k}{\{l : A\} \bowtie \{k : B\}} \hspace{3cm} \frac{A \bowtie B}{\{l : A\} \bowtie \{l : B\}}$$

The first states that any records with different labels are mergeable, because such types share no operations in common. The second states that even when two types both contain the field l, they may be merged if the contents of the field themselves are mergeable. In this case, some later part of every operation on the two types (after the projection of l) will disambiguate a use of this type.

These two rules follow directly from the first two rules of Figure 6. Functions can be merged when their arguments are distinguishable to form an overloaded function. Two functions that take the same argument type can also be merged when their outputs can be. We previously saw an example in which merging a constructor of type **Int** $\to$ {age : **Int**} with another of type **Int** $\to$ {$name$: **String**} yields a combined constructor of type **Int** $\to$ {age : **Int**, $name$: **String**} This style of merge supports nested composition and is also possible for universally quantified types. A merge of values of types A and A' (itself having type $A \sqcap A'$) can be merged with a third value of type B when both A and A' are mergeable with B. This ensures that operations on B do not conflict with those of A or A'. Since a value of type $A \sqcup A'$ may actually have either type A or A', for this union type to be mergeable with another type, both A and A' must be. A value of type $\top$ cannot be used in any way. Thus this type can trivially be merged with any other type. Finally, because application of terms to types is explicit in $F_{\bowtie}$, it is safe for function types and universal quantifiers to be mergeable with each other. In systems where universal quantifiers have no explicit elimination form, this may not be possible.

Distinguishability. Figure 6 also gives the distinguishability relation. Two types A and B are distinguishable when patterns of each type do not match values of the other type. In other words, distinguishability means that the pattern matching procedure can tell values of one type from values of the other. Most of the distinguishability rules are straightforward or precisely dual to those of mergeability. A key rule states that unequal type constructors are distinguishable. Another rule closes distinguishability over the subtyping relation. As shown in §6.2, this property is important for soundness of the type system.

Well-Formed Types. The well-formedness relation on types, written $\Gamma \vdash A$ ensures that:

(1) Only the type variables in Γ appear free in A.
(2) The components of every intersection in A are mergeable.
(3) The components of every union in A are distinguishable.

The cases for unions and intersection types are the only non-standard part of the definition of this relation, which is given in Figure 7.

The disjointness restrictions considerably simplify the type soundness argument by enabling inversion principles that we will see in §6. In the case of unions, the restrictions similarly simplify the pattern matching procedure. Consider matching the value $v = \lambda\{(\ast \hookrightarrow A \sqcup B) \mapsto v'\}$ against the pattern $p = (x : A)|(x : B)$. The matching procedure we saw earlier would try to match v against $x : A$ and $x : B$ separately. But these matches are potentially ill typed: $A \sqcup B$ is in general a subtype of neither A nor B. With the disjointness restriction on unions, it can be proven that every value of type $A \sqcup B$ either has type A or type B. Thus, one of the two matches is well typed.

A notable drawback of the disjointness restriction is that it precludes employing intersections as type refinements [Freeman and Pfenning 1991]. It appears that, under a notion of subtyping with the right distributivity rules, non-distinguishable unions (and their patterns) can often be simplified to distinguishable ones. Future work may study such an approach to support arbitrary union and intersection types in $F_{\bowtie}$. For our present purposes, disjoint intersections and unions capture the essence of what is necessary to investigate overloading and extensible data types.

4.3 Typing

Figure 8 defines $F_{\bowtie}$'s type system. It borrows some concepts from bidirectional typing which help ensure that the rules are type-directed. The judgment $\Gamma \vdash e \Rightarrow A$ ("e synthesizes type A under Γ") means that A is the least type which can be assigned to e. In other words, A is the principal type of e. Meanwhile, $\Gamma \vdash e \Leftarrow A$ ("e checks at type A under Γ") signifies that e has the type A and potentially some subtypes of A. It is defined by a single rule and holds precisely when e synthesizes some subtype of A under Γ. Unlike typical bidirectional systems, the distinction between checking and synthesis is not exploited to reduce the need for type annotations.

Applications are typed with the help of a type-level dispatch operator $\mathrm{disp}(A, \hat{B}) \Rightarrow C$. This is similar in style to the *apptype* function in the system of Freeman and Pfenning [1991]. The disp metafunction takes the type of a function A and a context $\hat{B}$ and computes the output type C. Due to overloading and union types, the implementation of this operation is non-trivial. It must statically determine which overloads of a function will execute. We will cover this in §5. It is, however, expected to abide by the following spec:

LEMMA 4.1 (SOUNDNESS AND COMPLETENESS OF TYPE-LEVEL DISPATCH).
(1) *We have* $\mathrm{disp}(A, B) \Rightarrow C$ *iff C is the least type such that $A \leq B \to C$.*
(2) *We have* $\mathrm{disp}(A, [B]) \Rightarrow C$ *iff C is the least type such that there exists A' where $A \leq \forall \alpha.A'$ and $C \leq A'[B/\alpha]$.*

Merges are typed by taking the intersection of the types they are annotated with. These types must be mergeable. Types are assigned to a λ-value by obtaining a type for each clause in its body, ensuring that the inferred types of all clauses are disjoint, and constructing an intersection of those types. This involves an auxiliary judgment $\Gamma; \hat{p} \vdash e \Rightarrow A$ which infers the type A of a single clause $\hat{p} \mapsto e$. Inferring the type of a clause involves determining the types of the bound variables in $\hat{p}$ as well as this pattern's return type annotation. These are used to type the expression e.

The next two judgments type patterns. Value patterns are typed with the judgment $\Gamma \vdash p \Rightarrow A \dashv \Delta$. The environment Γ contains the type variables free in p (as well as A, B, and Δ). The type A is the type of value that p matches (i.e. the type that p eliminates). Lastly, Δ contains the type and term variables bound by p.

An or-pattern binds only the variables that around bound by both subpatterns. This ensures that at runtime there is always a value to assign to every bound variable, even if only one subpattern matches. On the other hand, an and-pattern binds the variables that are mentioned on either side.

$$\boxed{\Gamma \vdash e \Leftrightarrow A} \hspace{4cm} \text{(expression type synthesis/checking)}$$

$$\frac{x : A \in \Gamma}{\Gamma \vdash x \Rightarrow A} \qquad \frac{\Gamma \vdash e \Rightarrow A}{\Gamma \vdash c\, e \Rightarrow c\, A} \qquad \frac{\Gamma \vdash e \Leftarrow A}{\Gamma \vdash (e : A) \Rightarrow A} \qquad \frac{\Gamma \vdash e \Rightarrow A \qquad A \leq B}{\Gamma \vdash e \Leftarrow B}$$

$$\frac{\Gamma \vdash e_2 \Rightarrow A_2}{\Gamma \vdash e_1 \Rightarrow A_1 \qquad \mathrm{disp}(A_1, A_2) \Rightarrow B}{\Gamma \vdash e_1\, e_2 \Rightarrow B} \qquad \frac{\Gamma \vdash A_2}{\Gamma \vdash e \Rightarrow A_1 \qquad \mathrm{disp}(A_1, [A_2]) \Rightarrow B}{\Gamma \vdash e\, [A_2] \Rightarrow B}$$

$$\frac{\Gamma \vdash A_1^- \quad \Gamma \vdash A_2^- \quad A_1^- \bowtie A_2^-}{\Gamma \vdash e_1 \Leftarrow A_1^- \quad \Gamma \vdash e_2 \Leftarrow A_2^-}{\Gamma \vdash (e_1 : A_1^-) \,\mathbin{\|}\, (e_2 : A_2^-) \Rightarrow A_1^- \sqcap A_2^-} \qquad \frac{\Gamma; \hat{p}_1 \vdash e_1 \Rightarrow A_1, \ldots, \Gamma; \hat{p}_n \vdash e_n \Rightarrow A_n}{A_1 \bowtie \ldots \bowtie A_n}{\Gamma \vdash \lambda\{\hat{p}_1 \mapsto e_1, \ldots, \hat{p}_n \mapsto e_n\} \Rightarrow A_1 \sqcap \ldots \sqcap A_n}$$

$$\boxed{\Gamma \vdash p \Rightarrow A \dashv \Delta} \hspace{4cm} \text{(value pattern typing)}$$

$$\frac{\Gamma \vdash A}{\Gamma \vdash (x : A) \Rightarrow A \dashv x : A} \qquad \frac{\Gamma \vdash p_1 \Rightarrow A_1 \dashv \Delta_1 \quad \Gamma \vdash p_2 \Rightarrow A_2 \dashv \Delta_2 \quad A_1 \blacklozenge A_2}{\Gamma \vdash p_1 | p_2 \Rightarrow A_1 \sqcup A_2 \dashv \Delta_1 \sqcup \Delta_2}$$

$$\frac{\Gamma \vdash p \Rightarrow A \dashv \Delta}{\Gamma \vdash c\, p \Rightarrow c\, A \dashv \Delta} \qquad \frac{\Gamma \vdash p_1 \Rightarrow A_1 \dashv \Delta_1 \quad \Gamma \vdash p_2 \Rightarrow A_2 \dashv \Delta_2 \quad A_1 \bowtie A_2 \quad p_1 \sim p_2}{\Gamma \vdash p_1 \& p_2 \Rightarrow A_1 \sqcap A_2 \dashv \Delta_1 \sqcap \Delta_2}$$

$$\boxed{\Gamma; \hat{p} \vdash e \Rightarrow B} \;\; (\lambda \text{ clause typing}) \;\; \boxed{\Gamma \vdash \hat{p} \Rightarrow A \hookrightarrow B \dashv \Delta} \hspace{2cm} \text{(elimination pattern typing)}$$

$$\frac{\Gamma \vdash \hat{p} \Rightarrow A \hookrightarrow B \dashv \Delta}{\Gamma, \Delta \vdash e \Leftarrow B}{\Gamma; \hat{p} \vdash e \Rightarrow A} \qquad \frac{\Gamma \vdash B^-}{\Gamma \vdash (* \hookrightarrow B^-) \Rightarrow B^- \hookrightarrow B^- \dashv \cdot}$$

$$\frac{\Gamma \vdash B \quad \Gamma \vdash p \Rightarrow A \dashv \Delta}{\Gamma \vdash (p \hookrightarrow B) \Rightarrow (A \to B) \hookrightarrow B \dashv \Delta} \qquad \frac{\Gamma, \alpha \vdash B}{\Gamma \vdash (\alpha \hookrightarrow B) \Rightarrow \forall \alpha.B \hookrightarrow B \dashv \alpha}$$

$$\boxed{p_1 \sim p_2} \hspace{6cm} \text{(value pattern consistency)}$$

$$\frac{}{(x : A) \sim (y : B)} \qquad \frac{p_1 \sim p_2}{c\, p_1 \sim c\, p_2} \qquad \frac{p_1 \sim p_2 \quad p_1' \sim p_2}{p_1 \& p_1' \sim p_2} \qquad \frac{p_1 \sim p_2 \quad p_1 \sim p_2'}{p_1 \sim p_2 \& p_2'}$$

$$\boxed{\Delta_1 \sqcup \Delta_2 \text{ and } \Delta_1 \sqcap \Delta_2}$$

$$x : A \in (\Delta_1 \sqcup \Delta_2) \quad \text{iff} \quad x : A_1 \in \Delta_1 \text{ and } x : A_2 \in \Delta_2 \text{ and } A = A_1 \sqcup A_2$$

$$\begin{aligned}
x : A_1 \in (\Delta_1 \sqcap \Delta_2) &\quad \text{if} \quad x : A_1 \in \Delta_1 \text{ and } x \notin \mathrm{dom}(\Delta_2) \\
x : A_2 \in (\Delta_1 \sqcap \Delta_2) &\quad \text{if} \quad x : A_2 \in \Delta_2 \text{ and } x \notin \mathrm{dom}(\Delta_1) \\
x : A_1 \sqcap A_2 \in (\Delta_1 \sqcap \Delta_2) &\quad \text{if} \quad x : A_1 \in \Delta_1 \text{ and } x : A_2 \in \Delta_2
\end{aligned}$$

Fig. 8. Typing Rules

The $\Delta_1 \sqcup \Delta_2$ and $\Delta_1 \sqcap \Delta_2$ operation take care of building the output environments in the typing of these kinds of patterns. Their definitions are reminiscent of classical record subtyping.

One subtle point in the definition of value pattern matching is required to ensure determinism. We need to rule out certain patterns like $(x : \top) \& c_0\, (x : (c_1\, \top \sqcup c_2\, \top))$ that bind the same variable twice to distinct values. This is achieved by requiring that the two sides of an and-pattern are *consistent*, written $p_1 \sim p_2$. Elimination patterns are typed with the judgment $\Gamma \vdash \hat{p} \Rightarrow A \hookrightarrow B \dashv \Delta$. The environments Γ and Δ perform the same functions as for value patterns. The type A, this time,

is the type of context that $\hat{p}$ matches (i.e. the type that $\hat{p}$ introduces). Since elimination patterns contain output type annotations, B is type of the expression bound under $\hat{p}$.

Example. Consider the value $\lambda\{(c\ (f : \textbf{Bool} \to \textbf{Bool}) \hookrightarrow \textbf{Int}) \mapsto \textbf{if}\ f\ \textbf{true}\ \textbf{then}\ 1\ \textbf{else}\ 0\}$ of type $(c\ (\textbf{Bool} \to \textbf{Bool})) \to \textbf{Int}$. This function pattern matches on its input to extract a function f from underneath a constructor c. Then it calls the function and returns an integer based on the result.

The elimination pattern $\hat{p} = c\ (f : \textbf{Bool} \to \textbf{Bool}) \hookrightarrow \textbf{Int}$ is typed with the following derivation:

$$\frac{\dfrac{}{\cdot \vdash (f : \textbf{Bool} \to \textbf{Bool}) \Rightarrow \textbf{Bool} \to \textbf{Bool} \dashv f : \textbf{Bool} \to \textbf{Bool}}}{\dfrac{\cdot \vdash c\ (f : \textbf{Bool} \to \textbf{Bool}) \Rightarrow c\ (\textbf{Bool} \to \textbf{Bool}) \dashv f : \textbf{Bool} \to \textbf{Bool}}{\cdot \vdash \hat{p} \Rightarrow (c\ (\textbf{Bool} \to \textbf{Bool}) \to \textbf{Int}) \hookrightarrow \textbf{Int} \dashv f : \textbf{Bool} \to \textbf{Bool}}}$$

At the bottom, in the elimination pattern judgment, note that the type **Int** appears twice. The first occurrence is to the right of the $\to$ in the type, which becomes the type of the entire function. The second is to the right of the $\hookrightarrow$ as the return type—that is, the type of the body **if** f **true then** 1 **else** 0.

5 TYPING ALGORITHM

The distinguishability, subtyping, and type-level dispatch relations have thus far been presented only in a declarative manner. We here discuss the algorithms required to compute these relations. As both of these play crucial roles in the type system, their implementations are necessary to show the typing relation is computable.

5.1 Splitting, Subtyping, and Disjointness

The main difficulty in our subtyping algorithm design arises from the distributivity over intersections and unions. A conventional technique adapted by many of the systems that support both intersection and union types is to convert types to standard forms before analyzing them [Castagna 2022; Freeman and Pfenning 1991]. We instead follow the idea of *splitting types* [Huang and Oliveira 2021], which normalizes types on the fly, and extend its subtyping to universally quantified and record types. The two splitting algorithms shown at the top of Figure 9 each take a type and transform it into an equivalent union or intersection type, respectively. The type $C \to A \sqcap B$, for instance, is equivalent to $(C \to A) \sqcap (C \to B)$, as we can derive $\text{split}_i(C \to A \sqcap B) \Rightarrow (C \to A) \sqcap (C \to B)$. We use a negated arrow to express that a type is not splittable, e.g., $\text{split}_u(\textbf{Int}) \not\Rightarrow$.

These rules encode all of the distributivity in subtyping and therefore take the burden off the algorithmic subtyping rules (defined in the middle of Figure 9). Once we replace these type splitting judgments by equations, i.e., write $A = A_1 \sqcap A_2$ and $B = B_1 \sqcup B_2$ instead of $\text{split}_i(A) \Rightarrow A_1 \sqcap A_2$ and $\text{split}_u(B) \Rightarrow B_1 \sqcup B_2$, the distributivity is fully eliminated, and we can see that the algorithmic subtyping rules follow a standard presentation. Importantly, it coincides with declarative subtyping.

LEMMA 5.1 (SOUNDNESS AND COMPLETENESS OF ALGORITHMIC SUBTYPING).
$A \le B$ *if and only if* $A \le_{\text{alg}} B$.

PROOF. Our proof follows that of Huang and Oliveira [2021]. The only if direction (soundness) is straightforward: it follows from the soundness of the type splitting relations.

The key to proving the if direction (completeness) is demonstrating the transitivity of the algorithmic system. For this, we induct on the sum of the sizes of the two types. When a type is split, we have an induction hypothesis for each piece. From there, the proof makes use of auxiliary inversion properties. □

$\boxed{\mathrm{split_u}(\hat{A}) \Rightarrow A_1 \sqcup A_2}$ $\hspace{5cm}$ *(union splitting)*

$$\frac{}{\mathrm{split_u}(A \sqcup B) \Rightarrow A \sqcup B} \qquad \frac{\mathrm{split_u}(A) \Rightarrow A_1 \sqcup A_2}{\mathrm{split_u}(\forall \alpha.A) \Rightarrow (\forall \alpha.A_1) \sqcup (\forall \alpha.A_2)} \qquad \frac{\mathrm{split_u}(A) \Rightarrow A_1 \sqcup A_2}{\mathrm{split_u}(c\, A) \Rightarrow c\, A_1 \sqcup c\, A_2}$$

$$\frac{\mathrm{split_u}(A) \Rightarrow A_1 \sqcup A_2}{\mathrm{split_u}(A \sqcap B) \Rightarrow (A_1 \sqcap B) \sqcup (A_2 \sqcap B)} \qquad \frac{\mathrm{split_u}(A) \nRightarrow \qquad \mathrm{split_u}(B) \Rightarrow B_1 \sqcup B_2}{\mathrm{split_u}(A \sqcap B) \Rightarrow (A \sqcap B_1) \sqcup (A \sqcap B_2)}$$

$\boxed{\mathrm{split_i}(A) \Rightarrow A_1 \sqcap A_2}$ $\hspace{5cm}$ *(intersection splitting)*

$$\frac{}{\mathrm{split_i}(A \sqcap B) \Rightarrow A \sqcap B} \qquad \frac{\mathrm{split_i}(A) \Rightarrow A_1 \sqcap A_2}{\mathrm{split_i}(\forall \alpha.A) \Rightarrow (\forall \alpha.A_1) \sqcap (\forall \alpha.A_2)} \qquad \frac{\mathrm{split_i}(A) \Rightarrow A_1 \sqcap A_2}{\mathrm{split_i}(c\, A) \Rightarrow c\, A_1 \sqcap c\, A_2}$$

$$\frac{\mathrm{split_i}(B) \Rightarrow B_1 \sqcap B_2}{\mathrm{split_i}(A \rightarrow B) \Rightarrow (A \rightarrow B_1) \sqcap (A \rightarrow B_2)} \qquad \frac{\mathrm{split_i}(B) \nRightarrow \qquad \mathrm{split_u}(A) \Rightarrow A_1 \sqcup A_2}{\mathrm{split_i}(A \rightarrow B) \Rightarrow (A_1 \rightarrow B) \sqcap (A_2 \rightarrow B)}$$

$$\frac{\mathrm{split_i}(A) \Rightarrow A_1 \sqcap A_2}{\mathrm{split_i}(A \sqcup B) \Rightarrow (A_1 \sqcup B) \sqcap (A_2 \sqcup B)} \qquad \frac{\mathrm{split_i}(A) \nRightarrow \qquad \mathrm{split_i}(B) \Rightarrow B_1 \sqcap B_2}{\mathrm{split_i}(A \sqcup B) \Rightarrow (A \sqcup B_1) \sqcap (A \sqcup B_2)}$$

$\boxed{A \leq_{\mathrm{alg}} B}$ $\hspace{5cm}$ *(algorithmic subtyping)*

$$\frac{}{A \leq_{\mathrm{alg}} A} \qquad \frac{A \leq_{\mathrm{alg}} B}{c\, A \leq_{\mathrm{alg}} c\, B} \qquad \frac{B_1 \leq_{\mathrm{alg}} A_1 \qquad A_2 \leq_{\mathrm{alg}} B_2}{A_1 \rightarrow A_2 \leq_{\mathrm{alg}} B_1 \rightarrow B_2} \qquad \frac{A \leq_{\mathrm{alg}} B}{\forall \alpha.A \leq_{\mathrm{alg}} \forall \alpha.B} \qquad \frac{}{\bot \leq_{\mathrm{alg}} A} \qquad \frac{}{A \leq_{\mathrm{alg}} \top}$$

$$\frac{\mathrm{split_i}(A) \Rightarrow A_1 \sqcap A_2 \quad A_1 \leq_{\mathrm{alg}} B}{A \leq_{\mathrm{alg}} B} \qquad \frac{\mathrm{split_i}(A) \Rightarrow A_1 \sqcap A_2 \quad A_2 \leq_{\mathrm{alg}} B}{A \leq_{\mathrm{alg}} B} \qquad \frac{\mathrm{split_i}(B) \Rightarrow B_1 \sqcap B_2 \quad A \leq_{\mathrm{alg}} B_1 \quad A \leq_{\mathrm{alg}} B_2}{A \leq_{\mathrm{alg}} B} \qquad \frac{\mathrm{split_u}(A) \Rightarrow A_1 \sqcup A_2 \quad A_1 \leq_{\mathrm{alg}} B \quad A_2 \leq_{\mathrm{alg}} B}{A \leq_{\mathrm{alg}} B}$$

$$\frac{\mathrm{split_u}(B) \Rightarrow B_1 \sqcup B_2 \qquad A \leq_{\mathrm{alg}} B_1}{A \leq_{\mathrm{alg}} B} \qquad \frac{\mathrm{split_u}(B) \Rightarrow B_1 \sqcup B_2 \qquad A \leq_{\mathrm{alg}} B_2}{A \leq_{\mathrm{alg}} B}$$

$\boxed{\mathrm{disp}(A, \hat{B}) \Rightarrow C}$ $\hspace{5cm}$ *(type-level dispatch)*

$$\frac{\mathrm{split_u}(B) \Rightarrow B_1 \sqcup B_2 \\ \mathrm{disp}(A, B_1) \Rightarrow C_1 \qquad \mathrm{disp}(A, B_2) \Rightarrow C_2}{\mathrm{disp}(A, B) \Rightarrow C_1 \sqcup C_2} \qquad \frac{\mathrm{split_u}(\hat{B}) \nRightarrow}{\mathrm{disp}(\bot, \hat{B}) \Rightarrow \bot} \qquad \frac{\mathrm{split_u}(B) \nRightarrow \qquad B \leq A}{\mathrm{disp}(A \rightarrow A', B) \Rightarrow A'}$$

$$\frac{}{\mathrm{disp}(\forall \alpha.A, [B]) \Rightarrow A[B/\alpha]} \qquad \frac{\mathrm{split_u}(\hat{B}) \nRightarrow \\ \mathrm{disp}(A_1, \hat{B}) \Rightarrow C_1 \qquad \mathrm{disp}(A_2, \hat{B}) \Rightarrow C_2}{\mathrm{disp}(A_1 \sqcup A_2, \hat{B}) \Rightarrow C_1 \sqcup C_2} \qquad \frac{\mathrm{split_u}(\hat{B}) \nRightarrow \\ \mathrm{disp}(A_1, \hat{B}) \Rightarrow C_1 \qquad \mathrm{disp}(A_2, \hat{B}) \nRightarrow}{\mathrm{disp}(A_1 \sqcap A_2, \hat{B}) \Rightarrow C_1}$$

$$\frac{\mathrm{split_u}(\hat{B}) \nRightarrow \\ \mathrm{disp}(A_1, \hat{B}) \nRightarrow \qquad \mathrm{disp}(A_2, \hat{B}) \Rightarrow C_2}{\mathrm{disp}(A_1 \sqcap A_2, \hat{B}) \Rightarrow C_2} \qquad \frac{\mathrm{split_u}(\hat{B}) \nRightarrow \\ \mathrm{disp}(A_1, \hat{B}) \Rightarrow C_1 \qquad \mathrm{disp}(A_2, \hat{B}) \Rightarrow C_2}{\mathrm{disp}(A_1 \sqcap A_2, \hat{B}) \Rightarrow C_1 \sqcap C_2}$$

Fig. 9. Subtyping and Related Algorithms

Since the distinguishability relation is closed over subtyping, deciding it also requires dealing with distributivity. Thus, the disjointness algorithm [Rioux et al. 2022b] uses splitting just as the subtyping algorithm does.

5.2 Typing Application

The meta function $\mathrm{disp}(A, B) \Rightarrow C$ is used to calculate the most precise output type for both applications and type applications. The function type A has a very general form as it describes overloaded functions or nested compositions. Metavariable $\hat{B}$ captures two cases: it is either the

term's type in an ordinary application, which we use B to represent, or the type argument in a type application, denoted by $[B]$. The definitions ensure that type arguments are not splittable. In other words, $\text{split}_u([B]) \Rightarrow$ holds trivially.

The base case for a type application is when a universal quantified type or a bottom type meets a type argument. Unlike in the function application rule, which checks subtyping, there is no side condition for type application, so we do not need to process the type argument. The dispatch function looks into all the universally quantified types in A, substitutes the type argument into their bodies and then re-composes them back into the original structure (all the bottom types are also kept in the process). For example, we have: $\text{disp}((\bot \sqcup \forall \alpha.A), [B]) \Rightarrow \bot \sqcup A[B/\alpha]$.

For an intersection type to be applicable, the dispatch function requires that at least one part of it is applicable. Imagine an overloaded function with two implementations typed by $A_1 \to A_2$ and $B_1 \to B_2$, respectively. The argument only needs to satisfy either A_1 or B_1 for at least one implementation to be applicable in the runtime. For unions, it is mandatory that both parts are applicable, as unions only guarantee that one side is satisfied.

Dispatch splits unions in the argument type eagerly. To see why, consider the following situation, where an intersection of two function types is applied to a union. The two function types can each take one of the possible argument types, but when directly compared with the whole argument type, neither can have the subtyping condition satisfied.

$$\text{disp}(((\mathbf{Int} \to A_1) \sqcap (\mathbf{Bool} \to A_2)), \mathbf{Int} \sqcup \mathbf{Bool}) \Rightarrow A_1 \sqcup A_2$$

After our dispatch function tears the argument type apart, both applications can proceed.

As the type of a term evolves during reduction, the dispatch function must be monotonic to ensure soundness. Specifically, the subtypes of two applicable types must also be applicable, and the return type computed by the dispatch function should be a subtype of the original return type.

LEMMA 5.2 (MONOTONICITY OF TYPE-LEVEL DISPATCH).
(1) *If* $\text{disp}(A, \hat{B}) \Rightarrow C$ *and* $A' \le A$ *then there exists* $C' \le C$ *such that* $\text{disp}(A', \hat{B}) \Rightarrow C'$.
(2) *If* $\text{disp}(A, B) \Rightarrow C$ *and* $B' \le B$ *then there exists* $C' \le C$ *such that* $\text{disp}(A, B') \Rightarrow C'$.

PROOF. Follows from the type-level dispatch specification given by Lemma 4.1. □

6 RESULTS

We prove the soundness of $F_{\bowtie}$'s type system using a progress and preservation argument [Wright and Felleisen 1994]. Given the presence of copatterns, our proof structure in some ways resembles that of Abel et al. [2013]. It requires a standard substitution lemma as well as inversion and canonical forms lemmas.

Thanks to the presence of elimination pattern matching, we need to reason not only about the shape of well-typed values, but also well-typed elimination frames. Figure 10 defines typing judgments for elimination frames. The type $\hat{A}$ of an elimination frame F is given by the judgment $\Gamma \vdash F \Rightarrow \hat{V}$. An alternative method of typing elimination frames is by thinking of them as contexts which, when filled with an expression of type A, produce an expression of type B. In this case, we write $\Gamma; A \vdash F \Rightarrow B$. With these definitions, we state our canonical forms lemma. This lemma describes the shape of a value or elimination frame given its type.

LEMMA 6.1 (CANONICAL FORMS).
(1) *If* $\cdot \vdash v \Leftarrow c\,A$ *then there exists* v' *such that* $v = c\,v'$.
(2) *If* $\cdot \vdash v \Leftarrow A \to B$ *then there exists* M *such that* $v = \lambda M$.
(3) *If* $\cdot \vdash v \Leftarrow \forall \alpha.A$ *then there exists* M *such that* $v = \lambda M$.
(4) *There is no* v *such that* $\cdot \vdash v \Leftarrow \bot$.

$$\textit{value types} \quad V \quad ::= \quad c\,V \mid A^-$$

$$\textit{elim. frame types} \quad \hat{V} \quad ::= \quad V \mid [A]$$

$\boxed{\Gamma \vdash F \Rightarrow \hat{V}}$ *(elim. frame typing)*

$\boxed{\Gamma; A \vdash F \Leftrightarrow B}$ *(elim. type synthesis/checking)*

$$\frac{\Gamma \vdash v \Rightarrow V}{\Gamma \vdash [\cdot]\,v \Rightarrow V} \qquad \frac{}{\Gamma \vdash [\cdot][A] \Rightarrow [A]}$$

$$\frac{\Gamma \vdash F \Rightarrow \hat{V} \quad \mathrm{disp}(A, \hat{V}) \Rightarrow B}{\Gamma; A \vdash F \Rightarrow B} \qquad \frac{\Gamma; A \vdash F \Rightarrow B' \quad B' \leq B}{\Gamma; A \vdash F \Leftarrow B}$$

$\boxed{V/A \Rightarrow^+ \checkmark}$ *(value coverage)*

$\boxed{\hat{V}/A^- \Rightarrow^- \checkmark}$ *(elimination coverage)*

$$\frac{V/A \Rightarrow^+ \checkmark}{c\,V/c\,A \Rightarrow^+ \checkmark} \qquad \frac{V/A \Rightarrow^+ \checkmark}{V/A \sqcup B \Rightarrow^+ \checkmark}$$

$$\frac{V/A \Rightarrow^+ \checkmark}{V/A \to B \Rightarrow^- \checkmark} \qquad \frac{}{[B]/\forall\,\alpha.A \Rightarrow^- \checkmark}$$

$$\frac{V/B \Rightarrow^+ \checkmark}{V/A \sqcup B \Rightarrow^+ \checkmark} \qquad \frac{V/A \Rightarrow^+ \checkmark \quad V/B \Rightarrow^+ \checkmark}{V/A \sqcap B \Rightarrow^+ \checkmark}$$

$$\frac{\hat{V}/A_1^- \Rightarrow^- \checkmark \quad \hat{V}/A_2^- \Rightarrow^- \checkmark}{\hat{V}/A_1^- \sqcup A_2^- \Rightarrow^- \checkmark} \qquad \frac{\hat{V}/A_1^- \Rightarrow^- \checkmark}{\hat{V}/A_1^- \sqcap A_2^- \Rightarrow^- \checkmark}$$

$$\frac{}{V/A^- \Rightarrow^+ \checkmark}$$

$$\frac{\hat{V}/A_2^- \Rightarrow^- \checkmark}{\hat{V}/A_1^- \sqcap A_2^- \Rightarrow^- \checkmark}$$
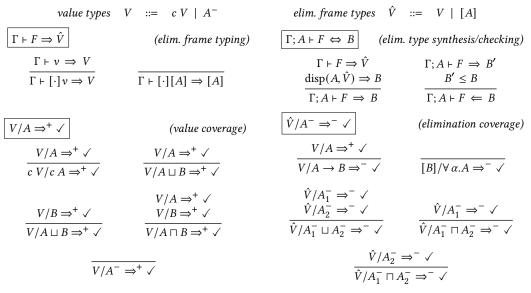
Fig. 10. Definitions Used in Soundness Proof

(5) *If $\cdot; A \to A' \vdash F \Rightarrow B$ then there exists v such that $F = [\cdot]\,v$.*
(6) *If $\cdot; \forall\,\alpha.A \vdash F \Rightarrow B$ then there exists C such that $F = [\cdot][C]$*
(7) *There are no F and B such that $\cdot; \top \vdash F \Rightarrow B$.*

PROOF. By case analysis on the value or frame in each part, using properties of subtyping. □

We also have a number of additional inversion principles for the typing relation. Most notable are those for union and intersection types. Perhaps surprisingly, inversion principles for unrestricted union types remain an active object of study [Castagna et al. 2022]. Given our disjointness restrictions, we are straightforwardly able to obtain:

LEMMA 6.2. *If $\Gamma \vdash v \Leftarrow A_1 \sqcup A_2$ and $A_1 \circledast A_2$ then $\Gamma \vdash v \Leftarrow A_1$ or $\Gamma \vdash v \Leftarrow A_2$.*

PROOF. By properties of subtyping. □

The distinguishability premise is crucial to this lemma, as is the fact that v is a value. Dually, we also have a property about frames that eliminate intersections of mergeable types.

LEMMA 6.3. *Suppose $\cdot \vdash A_1 \sqcap A_2$. If $\cdot; A_1 \sqcap A_2 \vdash F \Rightarrow B$ then there exists a B' such that either $\cdot; A_1 \vdash F \Rightarrow B'$ or $\cdot; A_2 \vdash F \Rightarrow B'$.*

PROOF. By properties of type-level dispatch. □

6.1 Progress

To prove progress, we must show that every closed, well-typed term is either a value or may take a step. The key *coverage* lemma needed here is that if a value v and a pattern p have type A, then v should successfully match p. An analogous property for elimination frames and elimination patterns is also needed. There are a couple of issues. First, A is not necessarily the principal type of v, so coverage is fundamentally a property relating subtyping to pattern matching. The subtyping relation is relatively complex, so proceeding directly by induction on it is tricky. Second, the type A does not uniquely determine the shape of the pattern p, which would require a direct proof to analyze many cases.

To deal with these issues, we define the value coverage relation $V/A \Rightarrow^+ \checkmark$, where V is the principal type of a value (also called a value type and described by Figure 10) and A is the type of a pattern. This relation holds when values of type V match patterns of type A; the defining rules of the relation in Figure 10 closely correspond to the rules defining the pattern matching procedure from Figure 4. To abstract over elimination pattern matching, we similarly define an elimination coverage relation $\hat{V}/A^- \Rightarrow^- \checkmark$. Here, $\hat{V}$ is the type of an elimination frame. The matching abstraction lemmas describe the intent behind both of these definitions. That is, if a value's (resp. elimination frame's) type matches a pattern's type, then the value (resp. elimination frame) matches the pattern.

LEMMA 6.4 (PATTERN MATCHING ABSTRACTION).
(1) *Suppose* $\cdot \vdash v \Rightarrow V$ *and* $\cdot \vdash A$. *If* $V/A \Rightarrow^+ \checkmark$ *then* v *matches all patterns* p *for which there exists some* Δ *such that* $\cdot \vdash p \Rightarrow A \dashv \Delta$.
(2) *Suppose* $\cdot \vdash F \Rightarrow \hat{V}$ *and* $\cdot \vdash \hat{p} \Rightarrow B^- \hookrightarrow C \dashv \Delta$. *If* $\hat{V}/B^- \Rightarrow^- \checkmark$ *then* F *matches* $\hat{p}$.

PROOF. For the first case, proceed by induction on $V/A \Rightarrow^+ \checkmark$ and $\hat{V}/B^- \Rightarrow^- \checkmark$ using inversion properties of pattern typing and subtyping. The proof of the second case is analogous. □

This lemma allows us to reason about pattern matching purely at the type level; we need not worry about the particular values or patterns involved. It is an approach with a clear connection to abstract interpretation [Cousot and Cousot 1977].

It remains to connect subtyping and the type-level dispatch operator to the coverage relations.

LEMMA 6.5 (COMPLETENESS OF COVERAGE RELATIONS).
(1) *If* $V \leq A$ *then* $V/A \Rightarrow^+ \checkmark$.
(2) *If there exists* C *such that* $\mathrm{disp}(A^-, \hat{V}) \Rightarrow C$ *then* $\hat{V}/A^- \Rightarrow^- \checkmark$.

PROOF. Both are proved via induction on the premise and use some auxiliary lemmas that construct the coverage relation judgments when one involved type is splittable. The proof of the second part makes use of the first when A^- is an arrow type. □

Coverage now follows easily from the previous two lemmas.

LEMMA 6.6 (COVERAGE).
(1) *Suppose* $\cdot \vdash A$. *If* $\Gamma \vdash v \Leftarrow A$ *and* $\Gamma \vdash p \Rightarrow A \dashv \Delta$ *then* v *matches* p.
(2) *If* $\cdot \vdash F \Rightarrow \hat{V}$ *and* $\mathrm{disp}(A^-, \hat{V}) \Rightarrow B_1$ *and* $\cdot \vdash \hat{p} \Rightarrow A^- \hookrightarrow B_2 \dashv \Delta$ *then* F *matches* $\hat{p}$.

PROOF. We prove the first case; the second case follows analogously from the same lemmas. By the definition of $\Gamma \vdash v \Leftarrow A$, we have $\Gamma \vdash v \Rightarrow V$ where $V \leq A$. By Lemma 6.5, we have $V/A \Rightarrow^+ \checkmark$. Applying Lemma 6.4 to this, we have that v matches p. □

The desired progress lemma is a consequence of coverage.

LEMMA 6.7 (PROGRESS). *If* $\cdot \vdash e \Rightarrow A$ *then* e *is a value or there exists* e' *such that* $e \longmapsto e'$.

PROOF. By induction on the typing derivation for e. The cases for elimination forms rely upon Lemma 6.6. □

6.2 Preservation

We face a rather subtle barrier to proving preservation. Consider the well-typed application of an overloaded function to an argument:

$$e = \lambda\{(x : \mathbf{Int}) \hookrightarrow \mathbf{Int} \mapsto e_1, (x : \mathbf{Bool}) \hookrightarrow \mathbf{Bool} \mapsto e_2\}\ v$$

Assuming $\cdot \vdash v \Rightarrow A$ and $x : \mathbf{Int} \vdash e_1 \Leftarrow \mathbf{Int}$ and $x : \mathbf{Bool} \vdash e_2 \Leftarrow \mathbf{Bool}$, we find that e has a type B such that $\mathrm{disp}((\mathbf{Int} \rightarrow \mathbf{Int}) \sqcap (\mathbf{Bool} \rightarrow \mathbf{Bool}), A) \Rightarrow B$. The expression e reduces to the expression e' below:

$$e' = \mathrm{disp}(\{\hat{p}_1 \mapsto e_1, \hat{p}_2 \mapsto e_2\}, [\cdot] v)$$

Establishing type preservation requires showing that the type of e' is a subtype of B. The problem arises in trying to determine what B is. The type-level dispatch operator first splits the type A. This makes things difficult since the only thing we know about A is that it is the type of an arbitrary value v. Intuitively, it seems that B should be the intersection of the output type of each overload that accepts B as an input. Unfortunately, that is not always true. Suppose $A = \mathbf{Int} \sqcup \mathbf{Bool}$. In this case, A is a subtype of neither $\mathbf{Int}$ nor $\mathbf{Bool}$. On the other hand, $B \equiv \mathbf{Int} \sqcup \mathbf{Bool}$. Thankfully, it turns out this case is impossible. There is no value whose principal type is $\mathbf{Int} \sqcup \mathbf{Bool}$; it is not a value type. To take advantage of this fact, we first need a couple of properties relevant to the dispatch procedure.

LEMMA 6.8 (DOWNWARD CLOSURE OF DISTINGUISHABILITY). *If $A \circledast B$ and $B' \leq B$ then $A \circledast B'$.*

PROOF. Immediate from the distinguishability rules.   □

LEMMA 6.9 (DISPATCH). *Suppose $\mathrm{split}_u(B) \Rightarrow$ and $\mathrm{split}_u(B') \Rightarrow$. If $A_1 \bowtie A_2$ and $\mathrm{disp}(A_1, B) \Rightarrow C_1$ and $\mathrm{disp}(A_1, B') \not\Rightarrow$ and $\mathrm{disp}(A_2, B') \Rightarrow C_2'$ then $B \circledast B'$.*

PROOF. By induction on $A_1 \bowtie A_2$, using Lemma 6.8.   □

From these we can prove the required properties of the disp operator.

LEMMA 6.10 (TYPE-LEVEL DISPATCH ON VALUE TYPES).
Suppose $\cdot \vdash A_1 \sqcap A_2$ and $\mathrm{disp}(A_1 \sqcap A_2, \hat{V}) \Rightarrow B$ and $\mathrm{disp}(A_1, \hat{V}) \Rightarrow B_1$.

(1) *If $\mathrm{disp}(A_2, \hat{V}) \Rightarrow B_2$ then $B_1 \sqcap B_2 \leq B$.*
(2) *If $\mathrm{disp}(A_2, \hat{V}) \not\Rightarrow$ then $B_1 \leq B$.*

PROOF. By scrutinizing whether $\hat{V}$ is splittable and applying inversion properties of the type-level dispatch operator.   □

Thanks to this lemma, we know that we can compute the output type of an application of an overloaded function to an unkown value from the output types of each individual overload. Lemma 5.2 allows the subtyping relationships in the conclusion of each of these cases to be strengthened to an equivalence.

With that difficulty taken care of, we next turn our attention to two key properties of pattern matching. The first ensures that patterns of distinguishable types are mutually exclusive while the second ensures that substitutions produced from pattern matching are well-typed.

LEMMA 6.11 (DISTINGUISHABLE PATTERNS). *Suppose $\Gamma \vdash p_1 \Rightarrow A_1 \dashv \Delta_1$ and $\Gamma \vdash p_2 \Rightarrow A_2 \dashv \Delta_2$ where $A_1 \circledast A_2$. There does not exist a closed, well-typed value that matches both p_1 and p_2.*

PROOF. By properties of the coverage relations and induction on $A_1 \circledast A_2$.   □

LEMMA 6.12 (ADEQUACY).

(1) *If $\cdot \vdash p \Rightarrow A \dashv \Delta$ and $\Gamma \vdash v \Leftarrow A$ and $v/p \Rightarrow \sigma$ then $\Gamma \vdash \sigma \Leftarrow \Delta$.*
(2) *Suppose $\cdot \vdash A$, and $\cdot \vdash \hat{p} \Rightarrow A \hookrightarrow B_1 \dashv \Delta$, and $\cdot; A \vdash F \Rightarrow B_2$, and $F/\hat{p} \Rightarrow \sigma \hookrightarrow B_3$. Then $\cdot; B_1 \vdash \sigma \Leftarrow \Delta \hookrightarrow B_3$ and $B_3 \leq B_2$.*

PROOF. By induction on the pattern matching rules using Lemma 6.11. In the case of value pattern matching against or-patterns, we again make use of Lemma 6.6.   □

The judgement $\cdot; B_1 \vdash \sigma \Leftarrow \Delta \hookrightarrow B_3$ in the conclusion above means, first, that each value in σ has a type given by Δ and, second, that an elimination frame in σ has input and output types B_1 and B_3 respectively. We refer the reader to the technical appendix [Rioux et al. 2022b] for its formal definition. With adequacy established, we can now prove the type preservation lemma.

LEMMA 6.13 (PRESERVATION). *If* $\cdot \vdash e \Leftarrow A$ *and* $e \longmapsto e'$ *then* $\cdot \vdash e' \Leftarrow A$.

PROOF. By induction on $e \longmapsto e'$, making use of Lemma 5.2, Lemma 6.10, and Lemma 6.12. □

6.3 Soundness and Determinism

With progress and preservation proven, our main result is now within reach.

THEOREM 6.14 (TYPE SOUNDNESS). *If* $\cdot \vdash e \Leftarrow A$ *and* $e \longmapsto^* e'$ *then either* e' *is a value or* e' *can take another step.*

PROOF. Follows from Lemma 6.7 and Lemma 6.13. Note that e synthesizes some subtype of A by the definition of the checking relation. Thus, fulfilling the premise of the progress lemma is not a problem. □

This establishes that well-typed programs do not "go wrong" [Milner 1978]. Furthermore, the type system also ensures the determinism of reduction.

LEMMA 6.15 (DETERMINISM OF EVALUATION).
(1) *If* $\Gamma \vdash p \Rightarrow A \dashv \Delta$ *and* $\Gamma \vdash v \Leftarrow A$ *and* $v/p \Rightarrow \sigma_1$ *and* $v/p \Rightarrow \sigma_2$ *then* $\sigma_1 = \sigma_2$.
(2) *If* $\Gamma \vdash \hat{p} \Rightarrow A \hookrightarrow B \dashv \Delta$ *and* $\Gamma \vdash F \Rightarrow \hat{V}$ *and* $F/\hat{p} \Rightarrow \sigma_1 \hookrightarrow B_1$ *and* $F/\hat{p} \Rightarrow \sigma_2 \hookrightarrow B_2$ *then* $\sigma_1 = \sigma_2$ *and* $B_1 = B_2$.
(3) *Suppose* $\cdot \vdash e \Rightarrow A$. *If* $e \longmapsto e_1$ *and* $e \longmapsto e_2$ *then* $e_1 = e_2$.

PROOF. The only source of nondeterminism in the definition of pattern matching (and reduction) is when p has the form $p_1 | p_2$ where v matches both p_1 and p_2. By inversion on the typing derivation for $p_1 | p_2$, the type of p_1 must be distinguishable from that of p_2. This contradicts Lemma 6.11. □

7 RELATED WORK

The Merge Operator and Disjoint Intersection Types. Oliveira et al. [2016] proposed the λ_i calculus, which only allows intersections of disjoint types. The goal of the disjointness restriction was to address the ambiguity in Dunfield's calculus and prove the coherence of the elaboration. Various extensions to λ_i, as well as relaxations to the type system were proposed afterwards. Bi et al. [2018] relaxed the disjointness restriction, requiring it only on merges and allowing the use of unrestricted intersections. To enable nested composition, they added a more powerful subtyping relation based on the well-known BCD subtyping [Barendregt et al. 1983] relation, which supports distributivity rules for intersections over other type constructs.

Huang et al. [2021] proposed a new approach to model the type-directed semantics of calculi with a merge operator, allowing for a direct proof of determinism of the operational semantics. Runtime implicit (up)casting replaces coercive subtyping. For example, being annotated by an intersection type $A \sqcap B$ means that a value will be cast by A and B respectively and merged together. $F_{\bowtie}$ also uses a type-directed semantics, but it does not employ a casting relation. Instead, our dynamic semantics is based on η-expansion. At runtime, merges are rewritten without duplication and other annotations are discarded. In addition, unlike all previous calculi with disjoint intersection types, $F_{\bowtie}$ supports overloading and union types by using two, more refined, disjointness relations (distinguishability and mergeability). The mergeability relation in $F_{\bowtie}$ is closely related

to the disjointness relation in the previous calculi. At the moment, $F_{\bowtie}$ does not yet support unrestricted intersections/unions and disjoint polymorphism, but we plan to study these extensions in the future.

Refinement Types. Intersection and union types are used in refinement type systems to increase the type-level expressiveness [Davies and Pfenning 2000; Freeman and Pfenning 1991]. Such systems either do not support overloading, as there is no merge, or use it only as a mechanism to annotate the term differently [Dunfield 2012]. These type systems face a similar problem: they must calculate the return type for application, where the function can have an intersection or union type. Due to the lack of overloading, the subtyping relation can be stronger, as we discussed in §4.1, but reduction is type irrelevant and often erases annotations.

Overloading, Semantic Subtyping, and CDuce. Castagna et al. [1995] studied a restricted form of the merge operator for overloading. In their calculus $\lambda\&$ only merges of functions are allowed and only one function is selected for β-reduction, so nested composition is not supported. Instead of disjointness, the criteria in $\lambda\&$ for types of merged functions is: for any argument that they can take, a best-matching type always exists. This restriction forbids merging a function of **Int** → **Int** → **Int** and a function of **Int** → **Bool** → **Bool**, but it allows overlapping among overloaded functions, which violates disjointness. Nonetheless adding more implementations to an overloaded function may lead to a different implementation being choosen with a best-match semantics, which could be unexpected. With disjointness this cannot happen, since adding a new overlapping implementation to an existing merge would result in a type error.

The *semantic subtyping* [Frisch et al. 2008] approach gives types a set-theoretic interpretation. Calculi with semantic subtyping are equipped with very expressive subtyping relations, containing unions, intersections, negation types and various distributivity rules. $F_{\bowtie}$ employs syntactic subtyping. While its subtyping relation is quite expressive, it does not support negation types. Calculi with semantic subtyping support expressive forms of overloading, which is typically resolved using a *best-match* semantics, as in Castagna et al.'s work. However, existing calculi with semantic subtyping do not support a merge operator or nested composition, key features for compositional programming [Zhang et al. 2021]. In contrast, $F_{\bowtie}$ is designed to support compositional programming features, and nested composition in particular. It integrates features of two previously separate types of calculi: those with disjoint intersection types and those with overloading.

CDuce uses semantic subtyping. Its core calculus [Xu 2013] has a type-case expression which has two branches and takes an expression, whose type, after evaluating to a value, determines which branch will be executed. Due to its special typing rule, the type case can encode overloaded functions typed by an intersection type, making it very similar to disjoint merges. But it has no disjointness constraint and only compares the value's type against a given type then behaves like a if-then-else expression. The cost is: functions in CoreCDuce are explicitly annotated. In contrast, our distinguishability relation intentionally avoids the need for comparing a function value to a function type in the runtime.

Elimination Constructs for Union Types. Unlike sum types (or tagged unions), untagged unions are not always equipped by an explicit elimination construct. When first introduced by MacQueen et al. [1984], the typing rule for union allows it to be eliminated under any context that can handle both possibilities of the union. Unfortunately, this rule breaks type preservation unless the β-reduction is performed in parallel [Barbanera et al. 1995]. A sound calculus with this typing rule must avoid evaluating the term of a union type multiple times after it substitutes the same variable. Two kinds of restrictions are employed in the literature: van Bakel et al. [2000] only allows values to be typed with union types, while Dunfield and Pfenning [2003] types subterms of union

type only when they occur in an evaluation context. The latter kind is a generalization of a more stringent approach employed by Pierce [1991b] that limits the context to be a union elimination form.

For the above calculi, the same piece of code is executed no matter what runtime type the term has. But practically, many languages supporting union types choose to have a type-dependent elimination construct: Igarashi and Nagira [2006] proposed a case analysis expression for Featherweight Java. XDuce has a pattern matching expression [Hosoya and Pierce 2003]. Generally speaking, the elimination constructs in those languages offer a first-match semantics, where cases can overlap and reordering the cases may change the semantics of the program. In $F_{\bowtie}$ union types are eliminated directly with applications, where the function part can, but does not have to be overloaded. Overloading functions, compared to type-case expression (typed by unions), provide more precise typing results as it can discriminate the argument type. For example, applying $(\textbf{Int} \rightarrow \textbf{Int}) \sqcap (\textbf{Bool} \rightarrow \textbf{Bool})$ to $\textbf{Int}$ gives us $\textbf{Int}$ while the corresponding type-case, which has type $(\textbf{Int} \sqcup \textbf{Bool}) \rightarrow (\textbf{Int} \sqcup \textbf{Bool})$, can only gives us $\textbf{Int} \sqcup \textbf{Bool}$. Note that the intersection type is a (proper) subtype of the function type with distributivity laws in subtyping. Based on matching styles, Rehman et al. [2022]'s work is the closest to ours. They proposed a union elimination construct (a switch-case expression) based on disjointness. Their construct is inspired by the Ceylon language [King 2013]. Their disjointness relation is closely related to our notion of distinguishability. Although their system supports intersection types, it does not have a merge operator and all values have a principal type that is not an intersection or union, which eliminates many of the technical issues at the cost of significantly less expressive power.

DATA AVAILABILITY STATEMENT

This article's artifact [Rioux et al. 2022a] contains a Coq development formalizing the type-level metatheory of this paper including subtyping, dispatch, and disjointness.

ACKNOWLEDGMENTS

REFERENCES

Andreas Abel, Brigitte Pientka, David Thibodeau, and Anton Setzer. 2013. Copatterns: Programming Infinite Structures by Observations. In *Proceedings of the 40th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (Rome, Italy) *(POPL '13)*. Association for Computing Machinery, New York, NY, USA, 27–38.

Jean-Marc Andreoli. 1992. Logic programming with focusing proofs in linear logic. *Journal of logic and computation* 2, 3 (1992), 297–347.

Franco Barbanera, Mariangiola Dezani-Ciancaglini, and Ugo de'Liguoro. 1995. Intersection and Union Types: Syntax and Semantics. *Information and Computation* 119, 2 (June 1995), 202–230.

Henk Barendregt, Mario Coppo, and Mariangiola Dezani-Ciancaglini. 1983. A filter lambda model and the completeness of type assignment. *The journal of symbolic logic* 48, 04 (1983), 931–940.

Xuan Bi and Bruno C. d. S. Oliveira. 2018. Typed First-Class Traits. In *European Conference on Object-Oriented Programming (ECOOP)*.

Xuan Bi, Bruno C. d. S. Oliveira, and Tom Schrijvers. 2018. The Essence of Nested Composition. In *European Conference on Object-Oriented Programming (ECOOP)*.

Xuan Bi, Ningning Xie, Bruno C. d. S. Oliveira, and Tom Schrijvers. 2019. Distributive Disjoint Polymorphism for Compositional Programming. In *European Symposium on Programming (ESOP)*.

Richard Bird and Oege de Moor. 1996. *The Algebra of Programming*. Prentice-Hall. http://www.cs.ox.ac.uk/publications/books/algebra/

Luca Cardelli and John C. Mitchell. 1990. Operations on Records. In *Proceedings of the Fifth International Conference on Mathematical Foundations of Programming Semantics* (New Orleans, Louisiana, USA). Springer-Verlag, Berlin, Heidelberg, 22–52.

J. Carette, O. Kiselyov, and C. Shan. 2009. Finally tagless, partially evaluated: Tagless staged interpreters for simpler typed languages. *Journal of Functional Programming* 19, 05 (2009), 509–543.

Giuseppe Castagna. 1997. Unifying Overloading and λ-Abstraction: $\Lambda_{\{\}}$. *Theor. Comput. Sci.* 176, 1–2 (apr 1997), 337–345.

Giuseppe Castagna. 2022. Covariance and Controvariance: a fresh look at an old issue (a primer in advanced type systems for learning functional programmers). *Logical Methods in Computer Science* Volume 16, Issue 1 (Feb. 2022).

Giuseppe Castagna, Giorgio Ghelli, and Giuseppe Longo. 1995. A calculus for overloaded functions with subtyping. *Information and Computation* 117, 1 (feb 1995), 115–135.

Giuseppe Castagna, Mickaël Laurent, Kim Nguyễn, and Matthew Lutze. 2022. On Type-Cases, Union Elimination, and Occurrence Typing. *Proc. ACM Program. Lang.* 6, POPL, Article 13 (Jan 2022), 31 pages.

Adriana B Compagnoni and Benjamin C Pierce. 1996. Higher-order intersection types and multiple inheritance. *Mathematical Structures in Computer Science (MSCS)* 6, 5 (1996), 469–501.

William R. Cook, Walter Hill, and Peter S. Canning. 1989. Inheritance is Not Subtyping. In *Proceedings of the 17th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '90)*. Association for Computing Machinery, 125–135.

Mario Coppo and Mariangiola Dezani-Ciancaglini. 1978. A new type assignment for λ-terms. *Archiv. Math. Logik* 19 (Jan 1978), 139–156.

Patrick Cousot and Radhia Cousot. 1977. Abstract Interpretation: A Unified Lattice Model for Static Analysis of Programs by Construction or Approximation of Fixpoints. In *Proceedings of the 4th ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages* (Los Angeles, California) *(POPL '77)*. Association for Computing Machinery, New York, NY, USA, 238–252.

Rowan Davies and Frank Pfenning. 2000. Intersection Types and Computational Effects. In *Proceedings of the Fifth ACM SIGPLAN International Conference on Functional Programming (ICFP '00)*. Association for Computing Machinery, New York, NY, USA, 198–208.

Jana Dunfield. 2012. Annotations for Intersection Typechecking. In *Proceedings of the Sixth Workshop on Intersection Types and Related Systems (EPTCS, Vol. 121)*, Stéphane Graham-Lengrand and Luca Paolini (Eds.). 35–47.

Jana Dunfield. 2014. Elaborating Intersection and Union Types. *J. Functional Programming* 24, 2–3 (2014), 133–165.

Jana Dunfield and Frank Pfenning. 2003. Type assignment for intersections and unions in call-by-value languages. In *International Conference on Foundations of Software Science and Computation Structures*. Springer, 250–266.

Erik Ernst. 2001. Family Polymorphism. In *European Conference on Object-Oriented Programming (ECOOP)*.

Tim Freeman and Frank Pfenning. 1991. Refinement types for ML. In *Proceedings of the ACM SIGPLAN 1991 conference on Programming language design and implementation*. 268–277.

Alain Frisch, Giuseppe Castagna, and Véronique Benzaken. 2008. Semantic Subtyping: Dealing Set-Theoretically with Function, Union, Intersection, and Negation Types. *J. ACM* 55, 4, Article 19 (Sep 2008), 64 pages.

Haruo Hosoya and Benjamin C Pierce. 2003. XDuce: A statically typed XML processing language. *ACM Transactions on Internet Technology (TOIT)* 3, 2 (2003), 117–148.

Xuejing Huang and Bruno C. d. S. Oliveira. 2021. Distributing intersection and union types with splits and duality (functional pearl). *Proceedings of the ACM on Programming Languages* 5, ICFP (2021), 1–24.

Xuejing Huang, Jinxu Zhao, and Bruno C. D. S. Oliveira. 2021. Taming the Merge Operator. *Journal of Functional Programming* 31 (2021).

Atsushi Igarashi and Hideshi Nagira. 2006. Union types for object-oriented programming. In *Proceedings of the 2006 ACM symposium on Applied computing*. 1435–1441.

Gavin King. 2013. The Ceylon language specification, version 1.0. https://ceylon-lang.org/documentation/1.0/spec/

David MacQueen, Gordon Plotkin, and Ravi Sethi. 1984. An ideal model for recursive polymorphic types. In *Proceedings of the 11th ACM SIGACT-SIGPLAN symposium on Principles of programming languages*. 165–174.

Robin Milner. 1978. A Theory of Type Polymorphism in Programming. *J. Comput. System Sci.* 17 (Aug. 1978), 348–375.

Martin Odersky, Philippe Altherr, Vincent Cremet, Burak Emir, Sebastian Maneth, Stéphane Micheloud, Nikolay Mihaylov, Michel Schinz, Erik Stenman, and Matthias Zenger. 2004. *An overview of the Scala programming language*. Technical Report. EPFL.

Bruno C. d. S. Oliveira and William R. Cook. 2012. Extensibility for the Masses. In *European Conference on Object-Oriented Programming (ECOOP)*.

Bruno C. d. S. Oliveira, Zhiyuan Shi, and João Alpuim. 2016. Disjoint intersection types. In *Proceedings of the 21st ACM SIGPLAN International Conference on Functional Programming* (New York, NY, USA, 2016-09-04) *(ICFP 2016)*. Association

for Computing Machinery, 364–377.

Dominic Orchard and Tom Schrijvers. 2010. Haskell Type Constraints Unleashed. In *Proceedings of the 10th International Conference on Functional and Logic Programming* (Sendai, Japan) *(FLOPS'10)*. Springer-Verlag, Berlin, Heidelberg, 56–71.

Benjamin C Pierce. 1991a. *Programming with intersection types and bounded polymorphism.* Ph.D. Dissertation. CMU-CS-91-205, Carnegie Mellon University.

Benjamin C Pierce. 1991b. *Programming with intersection types, union types, and polymorphism.* Technical Report.

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

Baber Rehman, Xuejing Huang, Ningning Xie, and Bruno C. d. S. Oliveira. 2022. Union Types with Disjoint Switches. In *36th European Conference on Object-Oriented Programming, ECOOP 2022, June 6-10, 2022, Berlin, Germany (LIPIcs, Vol. 222)*, Karim Ali and Jan Vitek (Eds.).

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

John C. Reynolds. 1997. *Design of the Programming Language FORSYTHE.* Birkhauser Boston Inc., USA, 173–233.

Nick Rioux, Xuejing Huang, Bruno C. d. S. Oliveira, and Steve Zdancewic. 2022a. A Bowtie for a Beast (Artifact). https://doi.org/10.5281/zenodo.7409103

Nick Rioux, Xuejing Huang, Bruno C. d. S. Oliveira, and Steve Zdancewic. 2022b. *A Bowtie for a Beast (Technical Appendix).* Technical Report. MS-CIS-22-02, University of Pennsylvania.

Tiark Rompf and Nada Amin. 2016. Type Soundness for Dependent Object Types (DOT). In *Proceedings of the 2016 ACM SIGPLAN International Conference on Object-Oriented Programming, Systems, Languages, and Applications* (Amsterdam, Netherlands) *(OOPSLA 2016)*. Association for Computing Machinery, 624–641.

Jeremy G. Siek and Walid Taha. 2006. Gradual typing for functional languages. In *Proceedings of the Scheme and Functional Programming Workshop (Scheme)*. ACM, 81–92.

Steffen van Bakel, Mariangiola Dezani-Ciancaglini, Ugo de'Liguoro, and Yoko Motohama. 2000. *The minimal relevant logic and the call-by-value lambda calculus.* Technical Report. TR-ARP-05-2000, The Australian National University.

Philip Wadler. 1998. The expression problem. *Java-genericity mailing list* (1998). https://homepages.inf.ed.ac.uk/wadler/papers/expression/expression.txt

Mitchell Wand. 1989. Type Inference for Record Concatenation and Multiple Inheritance. In *Symposium on Logic in Computer Science (LICS)*.

A.K. Wright and M. Felleisen. 1994. A Syntactic Approach to Type Soundness. *Inf. Comput.* 115, 1 (Nov 1994), 38–94.

Zhiwu Xu. 2013. *Parametric Polymorphism for XML Processing Languages*. Ph.D. Dissertation. Université Paris-Diderot-Paris VII.

Noam Zeilberger. 2008. Focusing and Higher-Order Abstract Syntax. In *Proceedings of the 35th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (San Francisco, California, USA) *(POPL '08)*. Association for Computing Machinery, New York, NY, USA, 359–369.

Weixin Zhang, Yaozhu Sun, and Bruno C. d. S. Oliveira. 2021. Compositional Programming. *ACM Transactions on Programming Languages and Systems* (April 2021).