

# Kind Inference for the FreeST Programming Language

Bernardo Almeida      Andreia Mordido

Vasco T. Vasconcelos

LASIGE, Faculdade de Ciências, Universidade de Lisboa, Portugal

{bpdalmeida,afmordido,vvasconcelos}@ciencias.ulisboa.pt

We present a kind inference algorithm for the FREEST programming language. The input to the algorithm is FREEST source code with (possibly part of) kind annotations replaced by kind variables. The algorithm infers concrete kinds for all kind variables. We ran the algorithm on the FREEST test suite by first replacing kind annotation on all type variables by fresh kind variables, and concluded that the algorithm correctly infers all kinds. Non surprisingly, we found out that programmers do not choose the most general kind in 20% of the cases.

## 1 Introduction

Software systems usually handle resources such as files and communication channels. The correct usage of such resources generally follows a protocol that describes valid patterns of interactions. For example a file should be opened and eventually closed, after which no read or write operations should ever be performed. The case for communication channels is similar: channels are opened, messages are exchanged, channels may eventually be closed, after which no more messages should be exchanged. Session types [6, 7, 15] allow expressing elaborate protocols (for files and channels, for example) guaranteeing that protocols are obeyed by programs.

FREEST [1, 2, 3] is a concurrent functional programming language based on System F where processes communicate via heterogeneously typed-channels governed by context-free session types [16]. Context-free session types allow describing protocols such as the serialization of arithmetic expressions. Consider the following datatype for arithmetic expressions.

```
1 data Exp = Lit Int | Plus Exp Exp | Times Exp Exp
```

An `Exp` is either a literal with an integer (Lit `Int`), a sum of two sub-expressions (Plus `Exp Exp`) or the product of two sub-expressions (Times `Exp Exp`). To serialise a value of type `Exp` we use a session type such as the following.

```
2 type ExpC =  $\oplus$ {LitC: !Int, PlusC: ExpC;ExpC, TimesC: ExpC;ExpC}
```

The abbreviation `ExpC` defines the type of a channel as seen from the point of view of the writer. A channel of type `ExpC` offers a set of options `LitC`, `PlusC` and `TimesC`. If the first option is chosen, an integer must be sent (`!Int`), while, in the others, two (sub-) expressions are expected to be sent.

Now, suppose that `serialise` is a function that serialises an `Exp` on a channel `ExpC`.

```
3 serialise : Exp  $\rightarrow$  ExpC; a  $\rightarrow$  a
```

The function expects a channel whose initial part is of type `ExpC` and then behaves as `a`: `serialise` is thus polymorphic on `a`. It consumes the front of the channel (of type `ExpC`) and returns the unused part of the channel (of type `a`).

As simple as it may seem, the above code is not valid in the current version of FREEST. The actual code requires further annotations allowing to distinguish functional from session types as well as linear from unrestricted types. The distinction is materialised by classifying types with kinds.

In FREEST kinds are composed of a multiplicity and a basic kind. Multiplicities control the number of times a value may be used: exactly once (linear, 1) or zero or more (unrestricted, \*). Basic kinds distinguish functional types (T) from session types (S). The reason why FREEST requires kinds lies on polymorphism. If  $! \mathbf{Int}; ? \mathbf{Int}$  is undoubtedly a session type and  $\mathbf{Int} \rightarrow \mathbf{Bool}$  a functional type, the same does not apply to the polymorphic variable  $a$ . Is it a session type or a functional type? The answer depends on the base kind of  $a$ : if S or then it is a session type, if T then it is a functional type. Kinds are thus necessary to decide whether the types such as  $a; ! \mathbf{Int}$  are well-formed.

The datatype defined in line 1 is currently written in annotated form as follows.

```
4 data Exp:*T = Lit Int | Plus Exp Exp | Times Exp Exp
```

The kind annotation  $*T$ , says that the datatype is functional. As for the multiplicity, we chose the unrestricted usage so that it may be used as often as required. Notwithstanding, one may declare  $\text{Exp}$  of kind  $1T$ , in which case `serialise` must become a linear function (of type  $\text{Exp} \rightarrow \text{ExpC}; a \rightarrow a$ ).

Expanding the abbreviation and annotating the datatype in line 2 we get the following type.

```
5 type ExpC:1S = rec a:1S .  $\oplus\{\text{LitC}: ! \mathbf{Int}, \text{PlusC}: a; a, \text{TimesC}: a; a\}$ 
```

$\text{ExpC}$  defines a recursive type that is well-formed when the kind of its body, the external choice ( $\oplus$ ), is a subkind of the kind for the recursion variable. In this case, the recursion variable  $\text{ExpC}$  is annotated with  $1S$ , given that its body is itself a linear session.

Finally, the function `serialise` is currently written as follows.

```
6 serialise :  $\forall a:1S . \text{Exp} \rightarrow \text{ExpC}; a \rightarrow a$ 
```

The polymorphic variable  $a$  stands for the continuation channel; it must be a linear session. Annotating  $a$  with the unrestricted session  $*S$  would dictate that it can only be instantiated with `Skip`, the only unrestricted session type.

Even if kinds are necessary in the underlying theory of the FREEST language, they clutter the code. The code in lines 1–3 is easier to understand and quicker to write; programmers need not fight the subtleties of each kind. Note that once kinds are inferred, the prenex occurrences of  $\forall$  can be omitted. The algorithm that we present in this paper annotates all type variables with their kinds, converting the code in lines 1–3 to that in lines 4–6.

The works more closely related to FREEST are Quill [9], Affe [13], Alms [17],  $F^\circ$  [8], FuSe<sup>{}</sup> [11] and Linear Haskell [4]. All these languages feature substructural type systems for dealing with linear, functional and affine types (in the case of Affe).

Quill [9] is a language with linear types and a syntax similar to that of Haskell. Quill features a novel design that combines linear and functional types. Contrarily to FREEST, Quill does not use kind mechanisms to distinguish between linear and functional types, instead it uses type predicates (or, qualified types) to reason about linearity. Furthermore, Quill does not support subkinding. Quill also has a type inference algorithm which was proven sound and complete. Affe [13] is an ML-like language with support to linear, affine and unrestricted types. Like Quill, Affe uses kinds and constrained types to distinguish between linear and affine types. Affe supports subkinding and it is equipped with full principal type inference. Like Affe, Alms [17] is an ML-like language but is based on System  $F_{<}^\omega$ , the higher-order polymorphic  $\lambda$ -calculus with subtyping. Alms supports affine and unrestricted types. It features a rich kind system with dependent kinds, unions, and intersections. Moreover, Alms supports ML modules, allows to expose unrestricted types as affine which gives flexibility to library programmers

$m ::= * \mid \mathbf{1} \mid \varphi$	Multiplicity
$\nu ::= \mathbf{S} \mid \mathbf{T}$	Prekind
$\kappa ::= m\nu \mid \chi$	Kind
$\sharp ::= ! \mid ?$	Polarity
$\star ::= \oplus \mid \&$	View
$(\cdot) ::= \{\cdot\} \mid \langle \cdot \rangle$	Record
$T ::= \text{Skip} \mid \text{End} \mid \sharp T \mid \star(\ell: T)_{\ell \in L} \mid T;T \mid ()_m$ $\mid T m \rightarrow T \mid (\ell: T)_{\ell \in L} \mid \forall a^\kappa. T \mid \mu a^\kappa. T \mid a$	Type
$e ::= ()_m \mid x \mid \lambda_m x: T. e \mid \Lambda a^\kappa. v \mid e e \mid \{\ell=e_\ell\}_{\ell \in L} \mid \text{let } \{\ell=x_\ell\}_{\ell \in L} = e \text{ in } e$ $\mid \ell e \mid \text{let } ()_m = e \text{ in } e \mid \text{case } e \text{ of } \{\ell \rightarrow x_\ell\}_{\ell \in L} \mid e[T] \mid \text{match } e \text{ with } \{\ell \rightarrow x_\ell\}_{\ell \in L}$	Expression


Figure 1: The syntax of kinds and types with support for kind inference

and it is equipped with local type inference.  $F^\circ$  [8] is an extension of System F that uses kinds to distinguish between linear and unrestricted types. Similarly to Affe and Alms, it supports subkinding. Similarly to FREEST, but unlike Affe,  $F^\circ$  does not support quantification over kinds. The work closest to FREEST in terms of context-free session types is FuSe<sup>{}^{\dagger}</sup> [11]. Padovani proposed an alternative formulation of context-free session types in which code and types are aligned via extra annotations, something we decided to avoid in FREEST. Linear Haskell [4] is a proposal to bring linear types to Haskell. In Linear Haskell functions  $T \rightarrow U$  and  $T \multimap U$  describe how the arguments of the function are used. The latter form, inspired by linear logic [5], uses the argument  $T$  exactly once. In FREEST, annotated arrows  $T \multimap U$  or  $T \mathbf{1} \rightarrow U$  describe how the function is used (unbounded usage or exactly once). FREEST kinding system differentiates session from functional types. It also classifies types according to their usage, linear or unrestricted. Other systems consider these notions separately (or only one of them). The ideas behind our inference algorithm are similar to Quill and Affe, but the details are quite different since we do not use type qualifiers to reason about linearity.

## 2 The Syntax of Kinds, Types and Expressions

This section briefly introduces the notions of kinds, types and expressions; we refer the interested reader to previous work for details [1]. FREEST relies on a base set for type variables (denoted by  $a, b, c$ ) and another for labels (denoted by  $k, \ell$ ). For the purpose of kind inference, we further use *multiplicity variables* (denoted by  $\varphi$ ) and *kind variables* (denoted by  $\chi$ ). The syntax of kinds, types and expressions is in fig. 1.

*Multiplicities* are used to indicate the number of times a value can be used. They are either unrestricted ( $*$ ), which denotes an arbitrary number of usages, linear ( $\mathbf{1}$ ), indicating precisely one usage, or a multiplicity variable ( $\varphi$ ). The kinding system relies on two base kinds:  $\mathbf{S}$  for session types and  $\mathbf{T}$  for arbitrary types. *Kinds* are either the combination of a base kind and a multiplicity or a kind variable  $\chi$ . Since a value of an unrestricted type may be used zero or more times, and one with a linear type must be used exactly once, it should be clear that an unrestricted value can be used where



a linear one is expected. Similarly, the interpretation of base kinds should be such that a session type ( $\ast\mathbf{S}$ ,  $\mathbf{1S}$ ) can be used in place of an arbitrary type ( $\mathbf{1T}$ ). The subkind relation for non variables (denoted  $\kappa <: \kappa$ ) forms a lattice, as exhibited in the diagram.

*Session types* include **Skip** indicating no communication, **End** representing channels ready to be closed, output ( $!T$ ) and input ( $?T$ ) messages, internal ( $\&\{\ell: T_\ell\}_{\ell \in L}$ ) and external choices ( $\oplus\{\ell: T_\ell\}_{\ell \in L}$ ) and sequential composition ( $T;U$ ). *Functional types* are composed of linear  $()_1$  and unrestricted unit types  $()_*$ , linear  $T \mathbf{1} \rightarrow U$  and unrestricted  $T \ast \rightarrow U$  functions, records  $\{\ell: T_\ell\}_{\ell \in L}$ , variants  $\langle \ell: T_\ell \rangle_{\ell \in L}$  and universal types  $\forall a^\kappa. T$ . Recursive types  $\mu a^\kappa. T$  are either session or functional depending on  $\kappa$ . Type variables  $a$  may refer to recursion variables in recursive types or to polymorphic variables in universal types. A function capturing in its body a free linear variable must itself be linear.

*Expressions* include variables  $x$ , term abstraction  $\lambda_m x: T. e$  and application  $e e$ , type abstraction  $\Lambda a^\kappa. v$  and application  $e[T]$ , record  $\{\ell = e_\ell\}_{\ell \in L}$  and record elimination  $\text{let } \{\ell = x_\ell\}_{\ell \in L} = e \text{ in } e$ , unit  $()_m$  and unit elimination  $\text{let } ()_m = e \text{ in } e$ , injection in a variant  $\ell e$  and variant elimination  $\text{case } e \text{ of } \{\ell \rightarrow x_\ell\}_{\ell \in L}$ . The expressions for channel operations include channel creation, **new**  $T$ , and branching on a choice, **match**  $e$  with  $\{\ell \rightarrow x_\ell\}_{\ell \in L}$ . The remaining operations on channels—namely **new**, **send**, **receive** and **select**  $\ell$ —are all understood as constants (pre-defined variables).

Given that our goal is to infer kind annotations, the reader may wonder why we allow them in the source language, namely in polymorphic types  $\forall a^\kappa. T$ , in recursive types  $\mu a^\kappa. T$  and in type abstractions  $\Lambda a^\kappa. v$ . Programmers may, if they so wish, provide kind annotations in the source code. Such annotations are passed to the algorithm. For those omitted, a fresh kind variable  $\chi$  is generated in its place.

### 3 Kind Inference

Our approach to kind inference follows the established two-step process, wherein the first gathers constraints and the second resolves the constraints. The constraint generation step produces constraints in two forms:  $\kappa <: \kappa$  and  $\varphi = \bigsqcup_{\ell \in L} \text{mult}(\kappa_\ell)$ . The first form represents subkinding constraints, while the second represents equalities between multiplicity variables and the least upper bound of a given set of multiplicities. To enhance readability, we use shorthand notation  $\varphi = \text{mult}(\kappa)$  for  $\varphi = \bigsqcup \text{mult}(\kappa)$  and use  $\sqcup$  in infix format for binary sets.

**Constraint Generation from Types** Kind and multiplicity constraints are captured by judgement  $\Delta_{\text{in}} \vdash T_{\text{in}} : \kappa_{\text{out}} \Rightarrow \mathcal{C}_{\text{out}}$ . The judgement states that type  $T$  has kind  $\kappa$  under kinding context  $\Delta$  (a map from type variables to kinds), producing constraint set  $\mathcal{C}$ . To clarify the distinction between input and output, we use the subscript **in** for parameters and **out** for results.

We explain a core subset of the constraint generation rules, those in fig. 2 (the complete set is in fig. 4). Rule CG-Var reads the kind for type variable  $a$  (recursive or polymorphic) from the kinding context, generating no additional restrictions. Rule CG-Rec governs recursive types which can either be session or functional. The kind of the recursion variable is copied to the kinding context when analysing type  $T$ . A constraint  $\kappa' <: \kappa$  is generated to ensure that the kind  $\kappa'$  of the body of the recursive type is a subkind of the kind  $\kappa$  of the recursion variable. Rule CG-Arrow, deals with functions  $T m \rightarrow U$ . It applies the algorithm recursively to  $T$  and  $U$ , and assigns the kind  $m\mathbf{T}$  to the function type, where  $m$  comes from the arrow annotation. Rule CG-Rcd builds kinds and constraints for all elements in the record. It generates a new fresh multiplicity variable  $\varphi$ . The result is kind  $\varphi\mathbf{T}$  and the constraint set is composed of the union of  $\mathcal{C}_\ell$  for all  $\ell \in L$  and a new constraint associating variable  $\varphi$  to the least upper bound of the multiplicities of  $\kappa_\ell$ . In order to ensure that  $\varphi$  gets the expected multiplicity, all elements

$$\boxed{\Delta_{\text{in}} \vdash T_{\text{in}} : \kappa_{\text{out}} \Rightarrow \mathcal{C}_{\text{out}}}$$

<p>CG-VAR</p> $\Delta, a : \kappa \vdash a : \kappa \Rightarrow \emptyset$	<p>CG-REC</p> $\frac{\Delta, a : \kappa \vdash T : \kappa' \Rightarrow \mathcal{C}}{\Delta \vdash \mu a^{\kappa}. T : \kappa' \Rightarrow \mathcal{C} \cup \{\kappa' <: \kappa\}}$	<p>CG-ARROW</p> $\frac{\Delta \vdash T : \kappa_1 \Rightarrow \mathcal{C}_1 \quad \Delta \vdash U : \kappa_2 \Rightarrow \mathcal{C}_2}{\Delta \vdash T m \rightarrow U : m\mathbf{T} \Rightarrow \mathcal{C}_1 \cup \mathcal{C}_2}$
<p>CG-RCD</p> $\frac{\Delta \vdash T_{\ell} : \kappa_{\ell} \Rightarrow \mathcal{C}_{\ell} \quad \varphi \text{ fresh} \quad (\forall \ell \in L)}{\Delta \vdash \{\ell : T_{\ell}\}_{\ell \in L} : \varphi\mathbf{T} \Rightarrow \bigcup_{\ell \in L} \mathcal{C}_{\ell} \cup \{\varphi = \bigsqcup_{\ell \in L} \text{mult}(\kappa_{\ell}), \kappa_{\ell} <: \varphi\mathbf{T}\}}$	<p>CG-TABS</p> $\frac{\Delta, a : \kappa \vdash T : \kappa' \Rightarrow \mathcal{C} \quad \varphi \text{ fresh}}{\Delta \vdash \forall a^{\kappa}. T : \varphi\mathbf{T} \Rightarrow \mathcal{C} \cup \{\varphi = \text{mult}(\kappa')\}}$	

Figure 2: Constraint generation from types

must be subkinds of the kind of the record, that is  $\varphi\mathbf{T}$ . Thus, if at least one entry in the record is linear, then  $\varphi$  is also constrained to be linear. Rule CG-TABS adds the kind of the polymorphic variable to the typing context when checking the body  $T$ . It then assigns kind  $\varphi\mathbf{T}$  to the incoming type  $\forall a^{\kappa}. T$ , where the fresh multiplicity variable  $\varphi$  denotes the multiplicity of the kind of type  $T$ .

Type operator  $\text{mult}$  is fully resolved only after analysing expressions. At this point it can only be partially resolved. When applied to a kind of the form  $m\mathbf{V}$  operator  $\text{mult}$  rewrites into multiplicity  $m$ , that is,  $\text{mult}(m\mathbf{V}) = m$ .

As an example, let us consider the function that extracts the first element of a pair.

$$\text{fst} : \forall a^{\chi_a}. \forall b^{\chi_b}. \{\text{fst} : a, \text{snd} : b\} * \rightarrow a$$

The application of the rules in fig. 2, yields the constraint set  $\{\varphi_1 = \text{mult}(\varphi_2\mathbf{T}), \varphi_2 = \text{mult}(*\mathbf{T}), \varphi_3 = \text{mult}(\chi_a) \sqcup \text{mult}(\chi_b)\}$ . Solving the constraint set one obtains  $\{\varphi_1 = *, \varphi_2 = *, \varphi_3 = \text{mult}(\chi_a) \sqcup \text{mult}(\chi_b)\}$ . We resolve the indeterminacy of kind variables  $\chi_a$  and  $\chi_b$  by assuming that they both are  $\mathbf{1T}$ , the maximum of the kind lattice. The solution would then be  $\{\varphi_1 = *, \varphi_2 = *, \varphi_3 = \mathbf{1}, \chi_a = \mathbf{1T}, \chi_b = \mathbf{1T}\}$ .

We argue that assigning  $\mathbf{1T}$  (the maximum) to  $\chi_a$  and  $\chi_b$  is the preferred solution, since it is the less restrictive of all solutions. If we were to choose another kind, such as  $*\mathbf{T}$ , then it would be impossible to call function  $\text{fst}$  on linear values (of types with kind  $\mathbf{1T}$ ). We would, undesirably, be ruling out some perfectly well-behaved programs.

But is  $\mathbf{1T}$  the best kind for variables  $\chi_a$  and  $\chi_b$ ? The answer depends on the definition of  $\text{fst}$ .

$$\text{fst} = \Lambda a^{\chi_a}. \Lambda b^{\chi_b}. \lambda_* p : \{\text{fst} : a, \text{snd} : b\}. \text{let } \{\text{fst} = x, \text{snd} = y\} = p \text{ in } x$$

An examination of expression  $\text{let } \{\text{fst} = x, \text{snd} = y\} = p \text{ in } x$  reveals that the second element of the pair,  $y$ , is discarded. Hence,  $\chi_b$  must be unrestricted. Would  $\chi_b = \mathbf{1T}$  be chosen, then FREEST would complain about a linearity violation when type checking the function. In other words, constraint  $\chi_b <: *\mathbf{T}$  must be added to the constraint set, but an inspection of the type of  $\text{fst}$  alone does not provide enough information to generate such a constraint. In the following, we present rules that allow generating constraints such as  $\chi_b <: *\mathbf{T}$  by inspecting variable usage in expressions.

**Constraint Generation from Expressions** Constraints for expressions are derived from judgement  $\Delta_{\text{in}} \mid \Gamma_{\text{in}} \vdash e_{\text{in}} : T_{\text{out}} \Rightarrow \mathcal{C}_{\text{out}} \mid \Sigma_{\text{out}}$ . The judgement states that expression  $e$  has type  $T$  under kinding context  $\Delta$  and typing context  $\Gamma$ . It generates a constraint set  $\mathcal{C}$  and a usage context  $\Sigma$ . Typing contexts map term variables  $x$  to types  $T$ ; usage contexts map term variables  $x$  to the kind  $\kappa$  of their types. Usage

$$\boxed{\Delta_{\text{in}} \mid \Gamma_{\text{in}} \vdash e_{\text{in}} : T_{\text{out}} \Rightarrow \mathcal{C}_{\text{out}} \mid \Sigma_{\text{out}}}$$

$$\begin{array}{c}
\text{INF-VAR} \\
\frac{\Delta \vdash T : \kappa \Rightarrow \mathcal{C}}{\Delta \mid \Gamma, x : T \vdash x : T \Rightarrow \mathcal{C} \mid \{x : \kappa\}} \\
\text{INF-ABS} \\
\frac{\Delta \vdash T_1 : \kappa \Rightarrow \mathcal{C}_1 \quad \Delta \mid \Gamma, x : T_1 \vdash e : T_2 \Rightarrow \mathcal{C}_2 \mid \Sigma \quad \mathcal{C}_3 = \text{if isAbs } e \text{ then } \{\kappa <: m\mathbf{T}\} \text{ else } \emptyset}{\Delta \mid \Gamma \vdash \lambda_m x : T_1. e : T_1 m \rightarrow T_2 \Rightarrow \mathcal{C}_1 \cup \mathcal{C}_2 \cup \mathcal{C}_3 \cup \text{Weaken}(\Sigma, x, \kappa) \mid \Sigma \setminus x} \\
\text{INF-APP} \\
\frac{\Delta \mid \Gamma \vdash e_1 : T_1 m \rightarrow T_2 \Rightarrow \mathcal{C}_1 \mid \Sigma_1 \quad \Delta \mid \Gamma \vdash e_2 : T_1 \Rightarrow \mathcal{C}_2 \mid \Sigma_2 \quad \Delta \vdash T_1 m \rightarrow T_2 : \kappa \Rightarrow \mathcal{C}_3}{\Delta \mid \Gamma \vdash e_1 e_2 : T_2 \Rightarrow \mathcal{C}_1 \cup \mathcal{C}_2 \cup \mathcal{C}_3 \cup \text{Merge}(\Sigma_1, \Sigma_2) \mid \Sigma_1 \cup \Sigma_2} \\
\text{INF-RCDELIM} \\
\frac{\Delta \mid \Gamma \vdash e_1 : \{\ell : T_\ell\}_{\ell \in L} \Rightarrow \mathcal{C}_1 \mid \Sigma_1 \quad \Delta \mid \Gamma, (x_\ell : T_\ell)_{\ell \in L} \vdash e_2 : T \Rightarrow \mathcal{C}_2 \mid \Sigma_2 \quad \Delta \vdash T : \kappa \Rightarrow \mathcal{C}_3 \quad \Delta \vdash T_\ell : \kappa_\ell \Rightarrow \mathcal{C}_\ell \quad \mathcal{C} = \mathcal{C}_1 \cup \mathcal{C}_2 \cup \mathcal{C}_3 \cup \mathcal{C}_\ell \cup \text{Merge}(\Sigma_1, \Sigma_2) \cup \text{Weaken}(\Sigma_2, x_\ell, \kappa_\ell) \quad (\forall \ell \in L)}{\Delta \mid \Gamma \vdash \text{let } \{\ell = x_\ell\}_{\ell \in L} = e_1 \text{ in } e_2 : T \Rightarrow \mathcal{C} \mid (\Sigma_1 \cup \Sigma_2) \setminus \{x_\ell\}_{\ell \in L}}
\end{array}$$

Figure 3: Constraint generation from expressions

contexts enable reasoning about variable usage: if the variable is used exactly once, it may be linear, otherwise it must be unrestricted. Next, we define functions `Weaken` and `Merge`. The former checks whether variables are used in expressions. If a variable is not used, then the set with constraint  $\kappa <: *T$  is returned. The latter checks whether a variable is used more than once: if it appears in multiple usage contexts, it must also be unrestricted.

$$\text{Weaken}(\Sigma, x, \kappa) = \begin{cases} \emptyset & \text{if } x \in \Sigma \\ \{\kappa <: *T\} & \text{otherwise} \end{cases} \quad \text{Merge}(\Sigma_1, \Sigma_2) = \{\kappa <: *T \mid x : \kappa \in \Sigma_1 \cap \Sigma_2\}$$

We are now in a position to explain the rules for expressions, in fig. 3 (the complete set is in fig. 5). Rule `Inf-Var` is used to assign a type to a variable in a given typing context. The rule requires the type context  $\Gamma$  to contain an entry  $x : T$ . The constraints pertaining to type  $T$  are gathered in  $\mathcal{C}$ . To reflect the usage of  $x$ , the rule returns a singleton map  $x : \kappa$ , where  $\kappa$  is the kind of  $T$ . Rule `Inf-Abs` deals with abstractions  $\lambda_m x : T_1. e$ . It recursively calls the judgments on  $T_1$  and on  $e$  to collect constraint sets  $\mathcal{C}_1$ ,  $\mathcal{C}_2$  and usage context  $\Sigma$ . The rule uses a new predicate, `isAbs`  $e$ , which holds when  $e$  is an abstraction. Then, if  $e$  is a closure the kind of  $T_1$  must be a subkind of  $mT$ , where  $m$  is the multiplicity of the abstraction. This restriction ensures that unrestricted abstractions do not close over linear values. The result is type  $T_1 m \rightarrow T_2$  together with a constraint set composed of the union of  $\mathcal{C}_1$ ,  $\mathcal{C}_2$ ,  $\mathcal{C}_3$  and the result of `Weaken`. The `Weaken` function checks whether a variable is unused at the end of its scope. In this case, the lambda abstraction introduces term variable  $x$  and therefore, at the end of the scope, we have to check its usage. Rule `Inf-App` states that if  $e_1$  has type  $T_1 m \rightarrow T_2$  and  $e_2$  has type  $T_1$ , then the expression  $e_1 e_2$  has type  $T_2$ . The constraints  $\mathcal{C}$  and usage context  $\Sigma$  are computed by combining the results of the kind inference of  $e_1$ ,  $e_2$  and  $T$ . The final constraint set is the union of  $\Sigma_1$ ,  $\Sigma_2$ ,  $\Sigma_3$ , and the result of the `Merge` function which imposes that any variable found in both  $\Sigma_1$  and  $\Sigma_2$  must be unrestricted. The final usage context is  $\Sigma_1 \cup \Sigma_2$ . Rule `Inf-RcdElim` combines all previously discussed concepts: it evaluates expressions  $e_1$  and  $e_2$ , collecting  $\mathcal{C}_1, \mathcal{C}_2$  and  $\Sigma_1, \Sigma_2$ . The result is the type of  $e_2$ , a constraint set  $\mathcal{C}$ , which is the union of

$\mathcal{C}_1, \mathcal{C}_2, \mathcal{C}_3$ , the result of Merge on  $\Sigma_1$  and  $\Sigma_2$ , and the application of Weaken on  $\Sigma_2$  for all  $x_\ell: \kappa_\ell$  to check for unused variables. The resulting usage context is the combination of  $\Sigma_1$  and  $\Sigma_2$  with all entries for  $x_\ell$  removed.

When analysing constraint generation from the type for function `f st`, we intuitively concluded that the second element in the pair must be unrestricted because it is discarded. The application of rules in fig. 3, yield the constraint set  $\{\chi_b <: *T, \chi_a <: \varphi_1 T, \chi_b <: \varphi_1 T, \chi_a <: \varphi_0 T, \chi_b <: \varphi_0 T, \varphi_0 = \text{mult}(\chi_a) \sqcup \text{mult}(\chi_b), \varphi_1 = \text{mult}(\chi_a) \sqcup \text{mult}(\chi_b)\}$ . A solution for this set is  $\{\varphi_0 = \mathbf{1}, \varphi_1 = \mathbf{1}, \chi_a = \mathbf{1}T, \chi_b = *T\}$ . The kind variable  $\chi_b$  is set to  $*T$  as we predicted. The constraint set is computed by combining the constraint sets generated resulting from applying the judgement to all sub-expressions and the result of functions Merge and Weaken. First, we examine the Merge function: it takes contexts  $\{p: \kappa_p\}$  and  $\{x: \chi_a\}$  as input and calculates the intersection of the two contexts, adding a constraint  $\kappa <: *T$  for each element in the intersection. This process ensures that any variable that is used in both contexts is unrestricted. The Weaken function is used to verify if any newly introduced variable is eventually discarded. In our example, Weaken is applied to  $x: \chi_a$  and  $y: \chi_b$  against usage context  $\{p: \kappa_p, x: \chi_a\}$ . For  $y: \chi_b$  function Weaken proceeds as follows: since  $y$  is not present in the context, a new constraint  $\{\chi_b <: *T\}$  is added. On the other hand, since  $x$  is already in the context, no constraint is created.

**Constraint Solving** We now describe an algorithm to solve constraint sets.

1. Initialise all kind variables  $\chi$  to the maximum of the kind lattice,  $\mathbf{1}T$ . Likewise initialize all multiplicity variables  $\varphi$  to the maximum of multiplicities,  $\mathbf{1}$ . Store them in  $\sigma$ .
2. Iterate over each constraint in the set:
  - (a) If the constraint is of the form  $\chi <: \kappa$ , then update the entry for  $\chi$  in  $\sigma$  with the greatest lower bound of  $\kappa$  and  $\sigma(\chi)$ . For example, if  $\sigma = [\chi \mapsto \mathbf{1}T]$  and we are analysing constraint  $\chi <: *T$ , then the value for  $\chi$  in  $\sigma$  must be updated to  $\mathbf{1}T \sqcap *T = *T$ . After this step, we would have  $\sigma = [\chi \mapsto *T]$ .
  - (b) If the constraint is of the form  $\kappa <: \chi$ , then check whether  $\kappa$  and the kind for  $\chi$  in  $\sigma$  is in the subkind relation; if not then fail. For example, if  $\sigma = [\chi \mapsto \mathbf{1}T]$  and we are analysing constraint  $*T <: \chi$ , then we find that it is in the subkind relation since  $*T <: \mathbf{1}T$ . A failure would happen with  $\sigma = [\chi \mapsto \mathbf{1}T]$ .
  - (c) If the constraint is of the form  $\kappa_1 <: \kappa_2$  and neither of the elements is a kind variable, then check whether  $\kappa_1 <: \kappa_2$  is in the subkind relation; if not then fail. If not fail, then remove constraint  $\kappa_1 <: \kappa_2$  from the constraint set.
  - (d) If the constraint is a multiplicity constraint  $\varphi = \bigsqcup_{\ell \in L} \text{mult}(\kappa_\ell)$ , then compute the least upper bound of the multiplicities. If any  $\kappa_\ell$  is a kind variable ( $\chi$ ) or a base kind with a multiplicity variable ( $\varphi T$ ), we get its kind from  $\sigma$  (recall that all variables are in  $\sigma$  as per step 1). If the thus obtained kind is more restrictive than that for  $\varphi$  in  $\sigma$  (e.g.  $*$  against  $\sigma(\varphi) = \mathbf{1}$ ), then store it in  $\sigma$ . If  $\varphi = *$ , then remove the constraint from the set.
3. Repeat the process in step 2 until there are no further updates to be made.
4. If all constraints have been satisfied, then return the solution  $\sigma$ . Otherwise, the constraint set is unsatisfiable.

In the case of function `f st`, the constraints gathered by the rules in fig. 3 are as follows.

$$\begin{aligned} \chi_1 <: \varphi_0 T, \chi_0 <: \varphi_0 T, \chi_1 <: \varphi_1 T, \chi_0 <: \varphi_1 T, \chi_1 <: *T, \\ \varphi_0 = \text{mult}(\chi_0) \sqcup \text{mult}(\chi_1), \varphi_1 = \text{mult}(\chi_0) \sqcup \text{mult}(\chi_1) \end{aligned}$$

Category of annotation	Number of annotations in the source code	Number of more general annotations generated
Datatypes	129	0
Type abbreviations	206	7
Universal types	282	94
Explicit recursive types	23	10
Type abstractions	30	25
Total	670	136

Table 1: Distribution of annotations

We start with  $\sigma = [\chi_0 \mapsto \mathbf{1T}, \chi_1 \mapsto \mathbf{1T}, \varphi_0 \mapsto \mathbf{1}, \varphi_1 \mapsto \mathbf{1}]$ . Next we pick constraint  $\chi_1 <: \varphi_0\mathbf{T}$  and use item 2(a). We have,  $\chi_1 <: \mathbf{1T}$  since  $\sigma(\varphi_0) = \mathbf{1}$ . Given that  $\sigma(\chi_0)$  is equal to  $\mathbf{1T}$ , and subkinding is reflexive,  $\sigma(\varphi_0)$  remains as  $\mathbf{1T}$ . The process for the second constraint,  $\chi_1 <: \varphi_0\mathbf{T}$ , is similar. We analyse the constraint  $\chi_1 <: \mathbf{1T}$  since  $\sigma(\varphi_0) = \mathbf{1}$ . Also in this case item 2(a) does not change  $\sigma$ . The next two constraints,  $\chi_1 <: \varphi_1\mathbf{T}$  and  $\chi_0 <: \varphi_1\mathbf{T}$ , are also handled by item 2(a). Once again,  $\sigma$  is subject to no update. Now we pick constraint  $\chi_1 <: *\mathbf{T}$ . Under item 2(a) the algorithm computes the greatest lower bound of  $*\mathbf{T}$  and  $\mathbf{1T}$ , which is  $*\mathbf{T}$ , so  $\sigma$  is updated accordingly. For the last two constraints we use item 2(d). We read the values of  $\chi_0$  and  $\chi_1$  from  $\sigma$  and compute the least upper bound of  $\text{mult}(\mathbf{1T})$  and  $\text{mult}(*\mathbf{T})$  which yields  $\mathbf{1}$ . Both entries for  $\chi_0$  and  $\chi_1$  are already  $\mathbf{1}$  and therefore no update to  $\sigma$  is done. Since we analysed all constraints and  $\sigma$  was updated in this iteration of the algorithm, the fixed-point is not reached yet and so we go through each constraint once again. This time no update is made and therefore we terminate with  $\sigma = [\chi_0 \mapsto \mathbf{1T}, \chi_1 \mapsto *\mathbf{T}, \varphi_0 \mapsto \mathbf{1}, \varphi_1 \mapsto \mathbf{1}]$ .

The algorithm iteratively updates the values of the kind and multiplicity variables until no further updates can be made, that is, until a fixed point is reached. Since the kind lattice is finite, any sequence of updates must eventually converge to a fixed point. For the same reason, each constraint can only be updated a finite number of times. Therefore, the algorithm terminates after a finite number of iterations.

The running time of the constraint generation algorithm is linear on the size of the input expression; that of the constraint satisfaction algorithm is quadratic. In the worst case scenario the number of constraints is equal to the size of the expression. Each constraint can only update  $\sigma$  twice (when a more restrictive solution is found). The worst case happens when a different constraint performs an update in each iteration, forcing the algorithm to analyse all the constraints in each iteration. A sensible optimization removes the constraints from the constraint set also in items 2(a) and 2(b), after concluding that they cannot update  $\sigma$  to a more restrictive solution. Since the update can only be performed a constant number of times, the algorithm becomes linear on the size of the input expression.

**Evaluation** We implemented the algorithm and incorporated it in the FREEST interpreter. Then we conducted an evaluation to check the behaviour of the algorithm when used on FREEST source code. The evaluation consisted of replacing all the 670 kind annotations by fresh kind variables in the 232 valid programs in the FREEST test suite and standard library (total of 9131 lines of code), running the algorithm and checking whether the algorithm infers the annotations back.

Kind annotations are spread over datatypes, type abbreviations, universal types, recursive types, and type abstractions. The distribution of annotations is as in table 1. The small number of annotations in recursive types and type abstractions comes from the fact that they are usually introduced implicitly, either via type abbreviations (as in the code in line 2) or through compiler elaboration introducing type



abstractions  $\Delta a^k.v$  for functions accompanied by their signatures.

We concluded that the algorithm correctly inferred all annotations and found that 136 of the 670 annotations (that is, 20%) were too specific and could be relaxed to a more general kind. The largest number of more general annotations found by the algorithm come from universal types. We attribute this to the conservative nature of programmers: if we are developing Church encodings (heavy on polymorphism), why would one require linear type variables? The algorithm did not improve the kind for datatypes: datatypes are usually used in an unrestricted manner in programs. Moreover, in the test suite, they usually appear as the first argument (to be pattern-matched) of functions with unrestricted closures and therefore they cannot be linear.

For an example where the algorithm suggests a more general kind, consider function composition.

```
dot :  $\forall a:*T\ b:*T\ c:*T . (b \rightarrow c) \rightarrow (a \rightarrow b) \rightarrow a \rightarrow c$ 
dot f g x = f (g x)
```

If we only provide unrestricted arguments to dot, then there is no reason why the polymorphic variables a, b and c could not have kind \*T. However, we would be ruling out programs that apply dot to linear arguments. Consider the following program.

```
dot : (b  $\rightarrow$  c)  $\rightarrow$  (a  $\rightarrow$  b)  $\rightarrow$  a  $\rightarrow$  c
dot f g x = f (g x)

g : ?Int;End  $\rightarrow$  Int
g c = let (x, c) = receive c in close c ; x

main : Int
main =
  let (w, r) = new () in
  fork (\_ 1  $\rightarrow$  let w = send 5 w in close w);
  dot id g r
```

This program would be flagged as untypable because we instantiate the polymorphic variable a with the linear session type ?Int;End. Since there is no reason why a, b and c should be unrestricted, the algorithm assigns kind 1T to the three polymorphic variables.

## 4 Future Work

There are several avenues for future work. The most immediate is to prove the correctness of the algorithm with respect to the typing system. Then, equipped with kind inference, we may think of introducing a third base kind, that for session types that must be eventually closed (that reach type End). In this case we would require the kind of the argument to function new to be of the newly introduced kind. We further plan to study the possibility of quantifying over kinds or multiplicities for extra flexibility in programming.

**Acknowledgements** We thank the anonymous reviewers for their detailed comments that greatly contributed to improve the paper. This work was supported by FCT through project SafeSessions, ref. PTDC/CCI-COM/6453/2020, and the LASIGE Research Unit, ref. UIDB/00408/2020 and ref. UIDP/00408/2020.

## References

- [1] Bernardo Almeida, Andreia Mordido, Peter Thiemann & Vasco T. Vasconcelos (2022): *Polymorphic lambda calculus with context-free session types*. *Inf. Comput.* 289(Part), p. 104948, doi:10.1016/j.ic.2022.104948.
- [2] Bernardo Almeida, Andreia Mordido & Vasco T. Vasconcelos (2019): *FreeST, a concurrent programming language with context-free session types*. <https://freest-lang.github.io>. Last accessed 2023.
- [3] Bernardo Almeida, Andreia Mordido & Vasco T. Vasconcelos (2019): *FreeST: Context-free Session Types in a Functional Language*. In: *PLACES, EPTCS 291*, pp. 12–23, doi:10.4204/EPTCS.291.2.
- [4] Jean-Philippe Bernardy, Mathieu Boespflug, Ryan R. Newton, Simon Peyton Jones & Arnaud Spiwack (2018): *Linear Haskell: practical linearity in a higher-order polymorphic language*. *Proc. ACM Program. Lang.* 2(POPL), pp. 5:1–5:29, doi:10.1145/3158093.
- [5] Jean-Yves Girard (1987): *Linear Logic*. *Theor. Comput. Sci.* 50, pp. 1–102, doi:10.1016/0304-3975(87)90045-4.
- [6] Kohei Honda (1993): *Types for Dyadic Interaction*. In: *CONCUR, LNCS 715*, Springer, pp. 509–523, doi:10.1007/3-540-57208-2\_35.
- [7] Kohei Honda, Vasco Thudichum Vasconcelos & Makoto Kubo (1998): *Language Primitives and Type Discipline for Structured Communication-Based Programming*. In: *ESOP, LNCS 1381*, Springer, pp. 122–138, doi:10.1007/BFb0053567.
- [8] Karl Mazurak, Jianzhou Zhao & Steve Zdancewic (2010): *Lightweight linear types in system fdegree*. In Andrew Kennedy & Nick Benton, editors: *Proceedings of TLDI 2010: 2010 ACM SIGPLAN International Workshop on Types in Languages Design and Implementation, Madrid, Spain, January 23, 2010*, ACM, pp. 77–88, doi:10.1145/1708016.1708027.
- [9] J. Garrett Morris (2016): *The best of both worlds: linear functional programming without compromise*. In Jacques Garrigue, Gabriele Keller & Eijiro Sumii, editors: *Proceedings of the 21st ACM SIGPLAN International Conference on Functional Programming, ICFP 2016, Nara, Japan, September 18-22, 2016*, ACM, pp. 448–461, doi:10.1145/2951913.2951925.
- [10] Martin Odersky, Christoph Zenger & Matthias Zenger (2001): *Colored local type inference*. In Chris Hankin & Dave Schmidt, editors: *Conference Record of POPL 2001: The 28th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, London, UK, January 17-19, 2001*, ACM, pp. 41–53, doi:10.1145/360204.360207.
- [11] Luca Padovani (2019): *Context-Free Session Type Inference*. *ACM Trans. Program. Lang. Syst.* 41(2), pp. 9:1–9:37, doi:10.1145/3229062.
- [12] Benjamin C. Pierce & David N. Turner (2000): *Local type inference*. *ACM Trans. Program. Lang. Syst.* 22(1), pp. 1–44, doi:10.1145/345099.345100.
- [13] Gabriel Radanne, Hannes Saffrich & Peter Thiemann (2020): *Kindly bent to free us*. *Proc. ACM Program. Lang.* 4(ICFP), pp. 103:1–103:29, doi:10.1145/3408985.
- [14] John C. Reynolds (1974): *Towards a theory of type structure*. In Bernard J. Robinet, editor: *Programming Symposium, Proceedings Colloque sur la Programmation, Paris, France, April 9-11, 1974, Lecture Notes in Computer Science 19*, Springer, pp. 408–423, doi:10.1007/3-540-06859-7\_148.
- [15] Kaku Takeuchi, Kohei Honda & Makoto Kubo (1994): *An Interaction-based Language and its Typing System*. In: *PARLE, LNCS 817*, Springer, pp. 398–413, doi:10.1007/3-540-58184-7\_118.
- [16] Peter Thiemann & Vasco T. Vasconcelos (2016): *Context-free session types*. In Jacques Garrigue, Gabriele Keller & Eijiro Sumii, editors: *Proceedings of the 21st ACM SIGPLAN International Conference on Functional Programming, ICFP 2016, Nara, Japan, September 18-22, 2016*, ACM, pp. 462–475, doi:10.1145/2951913.2951926.

- [17] Jesse A. Tov & Riccardo Pucella (2011): *Practical affine types*. In Thomas Ball & Mooly Sagiv, editors: *Proceedings of the 38th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2011, Austin, TX, USA, January 26-28, 2011*, ACM, pp. 447–458, doi:10.1145/1926385.1926436.
- [18] J. B. Wells (1994): *Typability and Type-Checking in the Second-Order lambda-Calculus are Equivalent and Undecidable*. In: *LICS*, IEEE Computer Society, pp. 176–185, doi:10.1109/LICS.1994.316068.
- [19] Andrew K. Wright (1995): *Simple Imperative Polymorphism*. *LISP Symb. Comput.* 8(4), pp. 343–355.

$$\boxed{\Delta_{\text{in}} \vdash T_{\text{in}} : \kappa_{\text{out}} \Rightarrow \mathcal{C}_{\text{out}}}$$

CG-UNIT  $\Delta \vdash ()_m : m\mathbf{T} \Rightarrow \emptyset$     CG-VAR  $\Delta, a : \kappa \vdash a : \kappa \Rightarrow \emptyset$     CG-SKIP  $\Delta \vdash \text{Skip} : *S \Rightarrow \emptyset$     CG-END  $\Delta \vdash \text{End} : \mathbf{1}S \Rightarrow \emptyset$

CG-MSG  $\frac{\Delta \vdash T : \kappa \Rightarrow \mathcal{C}}{\Delta \vdash \#T : \mathbf{1}S \Rightarrow \mathcal{C}}$     CG-CH  $\frac{\Delta \vdash T_\ell : \kappa_\ell \Rightarrow \mathcal{C}_\ell \quad (\forall \ell \in L)}{\Delta \vdash \star(\ell : T_\ell)_{\ell \in L} : \mathbf{1}S \Rightarrow \bigcup_{\ell \in L} \mathcal{C}_\ell \cup \{\kappa_\ell <: \mathbf{1}S\}}$

CG-SEQ  $\frac{\Delta \vdash T : \kappa_1 \Rightarrow \mathcal{C}_1 \quad \Delta \vdash U : \kappa_2 \Rightarrow \mathcal{C}_2 \quad \varphi \text{ fresh}}{\Delta \vdash T;U : \varphi S \Rightarrow \mathcal{C}_1 \cup \mathcal{C}_2 \cup \{\kappa_1 <: \mathbf{1}S, \kappa_2 <: \mathbf{1}S, \varphi = \text{mult}(\kappa_1) \sqcup \text{mult}(\kappa_2)\}}$

CG-REC  $\frac{\Delta, a : \kappa \vdash T : \kappa' \Rightarrow \mathcal{C}}{\Delta \vdash \mu a^\kappa . T : \kappa' \Rightarrow \mathcal{C} \cup \{\kappa' <: \kappa\}}$     CG-ARROW  $\frac{\Delta \vdash T : \kappa_1 \Rightarrow \mathcal{C}_1 \quad \Delta \vdash U : \kappa_2 \Rightarrow \mathcal{C}_2}{\Delta \vdash T m \rightarrow U : m\mathbf{T} \Rightarrow \mathcal{C}_1 \cup \mathcal{C}_2}$

CG-RCD  $\frac{\Delta \vdash T_\ell : \kappa_\ell \Rightarrow \mathcal{C}_\ell \quad \varphi \text{ fresh} \quad (\forall \ell \in L)}{\Delta \vdash \{\ell : T_\ell\}_{\ell \in L} : \varphi \mathbf{T} \Rightarrow \bigcup_{\ell \in L} \mathcal{C}_\ell \cup \{\varphi = \bigsqcup_{\ell \in L} \text{mult}(\kappa_\ell), \kappa_\ell <: \varphi \mathbf{T}\}}$     CG-TABS  $\frac{\Delta, a : \kappa \vdash T : \kappa' \Rightarrow \mathcal{C} \quad \varphi \text{ fresh}}{\Delta \vdash \forall a^\kappa . T : \varphi \mathbf{T} \Rightarrow \mathcal{C} \cup \{\varphi = \text{mult}(\kappa')\}}$

Figure 4: Constraint generation from types (complete set of rules)

$$\boxed{\Delta_{\text{in}} \mid \Gamma_{\text{in}} \vdash e_{\text{in}} : T_{\text{out}} \Rightarrow \mathcal{C}_{\text{out}} \mid \Sigma_{\text{out}}}$$

$$\begin{array}{c}
\text{INF-CONST} \\
\frac{\Delta \vdash \text{typeof}(c) : \kappa \Rightarrow \mathcal{C}}{\Delta \mid \Gamma \vdash c : \text{typeof}(c) \Rightarrow \mathcal{C} \mid \emptyset} \\
\text{INF-VAR} \\
\frac{\Delta \vdash T : \kappa \Rightarrow \mathcal{C}}{\Delta \mid \Gamma, x : T \vdash x : T \Rightarrow \mathcal{C} \mid \{x : \kappa\}} \\
\text{INF-ABS} \\
\frac{\Delta \vdash T_1 : \kappa \Rightarrow \mathcal{C}_1 \quad \Delta \mid \Gamma, x : T_1 \vdash e : T_2 \Rightarrow \mathcal{C}_2 \mid \Sigma \quad \mathcal{C}_3 = \text{if isAbs } e \text{ then } \{\kappa <: mT\} \text{ else } \emptyset}{\Delta \mid \Gamma \vdash \lambda_m x : T_1. e : T_1 m \rightarrow T_2 \Rightarrow \mathcal{C}_1 \cup \mathcal{C}_2 \cup \mathcal{C}_3 \cup \text{Weaken}(\Sigma, x, \kappa) \mid \Sigma \setminus \{x : \kappa\}} \\
\text{INF-APP} \\
\frac{\Delta \mid \Gamma \vdash e_1 : T_1 m \rightarrow T_2 \Rightarrow \mathcal{C}_1 \mid \Sigma_1 \quad \Delta \mid \Gamma \vdash e_2 : T_1 \Rightarrow \mathcal{C}_2 \mid \Sigma_2 \quad \Delta \vdash T_1 m \rightarrow T_2 : \kappa \Rightarrow \mathcal{C}_3}{\Delta \mid \Gamma \vdash e_1 e_2 : T_2 \Rightarrow \mathcal{C}_1 \cup \mathcal{C}_2 \cup \mathcal{C}_3 \cup \text{Merge}(\Sigma_1, \Sigma_2) \mid \Sigma_1 \cup \Sigma_2} \\
\text{INF-TABS} \\
\frac{\Delta, a : \kappa \mid \Gamma \vdash v : T \Rightarrow \mathcal{C}_1 \mid \Sigma \quad \Delta \vdash T : \kappa' \Rightarrow \mathcal{C}_2}{\Delta \mid \Gamma \vdash \Lambda a^{\kappa}. v : \forall a^{\kappa}. T \Rightarrow \mathcal{C}_1 \cup \mathcal{C}_2 \mid \Sigma} \\
\text{INF-TAPP} \\
\frac{\Delta \vdash T : \kappa_1 \Rightarrow \mathcal{C}_1 \quad \Delta \mid \Gamma \vdash e : \forall a^{\kappa_2}. U \Rightarrow \mathcal{C}_2 \mid \Sigma}{\Delta \mid \Gamma \vdash e[T] : U[T/a] \Rightarrow \mathcal{C}_1 \cup \mathcal{C}_2 \mid \Sigma} \\
\text{INF-RCDELIM} \\
\frac{\Delta \mid \Gamma \vdash e_1 : \{\ell : T_\ell\}_{\ell \in L} \Rightarrow \mathcal{C}_1 \mid \Sigma_1 \quad \Delta \mid \Gamma, (x_\ell : T_\ell)_{\ell \in L} \vdash e_2 : T \Rightarrow \mathcal{C}_2 \mid \Sigma_2 \quad \Delta \vdash T : \kappa \Rightarrow \mathcal{C}_3}{\Delta \vdash T_\ell : \kappa_\ell \Rightarrow \mathcal{C}_\ell \quad \mathcal{C} = \mathcal{C}_1 \cup \mathcal{C}_2 \cup \mathcal{C}_3 \cup \text{Merge}(\Sigma_1, \Sigma_2) \cup \text{Weaken}(\Sigma_2, x_\ell, \kappa_\ell) \quad (\forall \ell \in L)} \\
\Delta \mid \Gamma \vdash \text{let } \{\ell = x_\ell\}_{\ell \in L} = e_1 \text{ in } e_2 : T \Rightarrow \mathcal{C} \mid (\Sigma_1 \cup \Sigma_2) \setminus \{x_\ell : \kappa_\ell\}_{\ell \in L} \\
\text{INF-RCD} \\
\frac{\Delta \mid \Gamma \vdash e_\ell : T_\ell \Rightarrow \mathcal{C}_\ell \mid \Sigma_\ell \quad \Delta \vdash T_\ell : \kappa_\ell \Rightarrow \mathcal{C}'_\ell \quad (\forall \ell \in L)}{\Delta \mid \Gamma \vdash \{\ell = v_\ell\}_{\ell \in L} : \{\ell : T_\ell\}_{\ell \in L} \Rightarrow \mathcal{C}_\ell \cup \mathcal{C}'_\ell \cup \text{Merge}(\Sigma_\ell) \mid \bigcup_{\ell \in L} \Sigma_\ell} \\
\text{INF-VARIANT} \\
\frac{\Delta \mid \Gamma \vdash e : T_k \Rightarrow \mathcal{C}_1 \mid \Sigma \quad \Delta \vdash T_\ell : \kappa_\ell \Rightarrow \mathcal{C}_\ell \quad k \in L \quad (\forall \ell \in L)}{\Delta \mid \Gamma \vdash k e : \langle \ell : T_\ell \rangle_{\ell \in L} \Rightarrow \mathcal{C}_1 \cup \mathcal{C}_\ell \mid \Sigma} \\
\text{INF-CASE} \\
\frac{\Delta \mid \Gamma \vdash e : \langle \ell : T_\ell \rangle_{\ell \in L} \Rightarrow \mathcal{C}_1 \mid \Sigma_1 \quad \Delta \mid \Gamma \vdash e_\ell : T_\ell m \rightarrow T \Rightarrow \mathcal{C}_\ell \mid \Sigma_\ell \quad \Delta \vdash T_\ell : \kappa_\ell \Rightarrow \mathcal{C}'_\ell \quad (\forall \ell \in L)}{\Delta \mid \Gamma \vdash \text{case } e \text{ of } \{\ell \rightarrow x_\ell\}_{\ell \in L} : T \Rightarrow \mathcal{C}_1 \cup \mathcal{C}_\ell \cup \mathcal{C}'_\ell \mid \Sigma_1 \cup \Sigma_\ell} \\
\text{INF-SEL} \\
\frac{\Delta \vdash T_\ell : \kappa_\ell \Rightarrow \mathcal{C}_\ell \quad \Delta \mid \Gamma \vdash e_\ell : T_\ell m \rightarrow T \Rightarrow \mathcal{C}'_\ell \mid \Sigma_\ell \quad k \in L \quad (\forall \ell \in L)}{\Delta \mid \Gamma \vdash \text{select } k : \oplus (\ell : T_\ell)_{\ell \in L} m \rightarrow T_k \Rightarrow \mathcal{C}_\ell \cup \mathcal{C}'_\ell \mid \bigcup_{\ell \in L} \Sigma_\ell} \\
\text{INF-NEW} \\
\frac{\emptyset \vdash T : \kappa \Rightarrow \mathcal{C}}{\Delta \mid \Gamma \vdash \text{new } T : \{\text{fst} : T, \text{snd} : \bar{T}\} \Rightarrow \mathcal{C} \mid \emptyset}
\end{array}$$

Figure 5: Constraint generation from expressions (complete set of rules)