



HAL
open science

From Amber to Coercion Constraints

Didier Rémy, Julien Cretin

► **To cite this version:**

Didier Rémy, Julien Cretin. From Amber to Coercion Constraints. Martin Abadi; Philippa Gardner; Andrew D. Gordon; Radu Mardare. Essays for the Luca Cardelli Fest, MSR-TR-2014-104, Microsoft Research, 2014, TechReport. hal-01093216

HAL Id: hal-01093216

<https://inria.hal.science/hal-01093216>

Submitted on 10 Dec 2014

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

From Amber to Coercion Constraints

Didier Rémy
INRIA

Julien Cretin
TrustInSoft

Abstract

Subtyping is a common tool in the design of type systems that finds its roots in the η -expansion of arrow types and the notion of type containment obtained by closing System F by η -expansion. Although strongly related, subtyping and type containment still significantly differ from one another when put into practice. We introduce *coercion constraints* to relate and generalize subtyping and type containment as well as all variants of F-bounded quantification and instance-bounded quantification used for first-order type inference in the presence of second-order types. We obtain a type system with a clearer separation between computational and erasable parts of terms.

1 The different flavors of subtyping

Subtyping in Amber Nowadays, subtyping is a well-understood concept and tool in the design of type systems, but it has not always been so. The origin of subtyping goes back to the 60's when type conversions between base types or type classes were introduced in programming languages, but the first formalization of subtyping is by Reynolds (1980). Subtyping was introduced in the strongly typed functional language **Amber** by Cardelli (1984). By contrast with simple type conversions between base types, subtyping in **Amber** can be propagated through arrow types covariantly on the codomains of functions and contravariantly on their domains.

The language **Amber** had only subtyping but no parametric polymorphism: its typing rules are those of the simply typed λ -calculus extended with a subtyping rule SUB:

$$\begin{array}{c} \text{SUB} \\ \frac{\Gamma \vdash a : \tau \quad \tau \triangleright \sigma}{\Gamma \vdash a : \sigma} \end{array} \qquad \begin{array}{c} \text{BOT} \\ \perp \triangleright \tau \end{array} \qquad \begin{array}{c} \text{TOP} \\ \tau \triangleright \top \end{array} \qquad \begin{array}{c} \text{ARROW} \\ \frac{\tau' \triangleright \tau \quad \sigma \triangleright \sigma'}{\tau \rightarrow \sigma \triangleright \tau' \rightarrow \sigma'} \end{array}$$

Through this paper we use the notation $\tau \triangleright \sigma$ to mean that τ is a subtype of σ (or, rather, τ *can be coerced to* σ) for homogeneity between different forms of coercions that we will encounter. The subtyping relation is defined by the three rules, BOT, TOP, and ARROW. For the sake of brevity we only consider the bottom and top types and arrow types; the language **Amber** also had records and subtyping between them which played a crucial role in the encoding of objects (but this is not our focus here), as well as recursive types which we introduce later on.

Type containment The notion of type containment was introduced simultaneously by Mitchell (1984, 1988) who considered the closure of System F by η -expansion, called F_η : by definition, a closed term is in F_η if it has an η -expansion in System F. Interestingly, F_η has also a syntactic presentation that is closely related to the notion of subtyping. It extends the typing rules of System F with a subtyping rule similar to, but richer than, the one of **Amber**. Indeed, System F has polymorphic types, which we write $\forall(\alpha) \tau$, and therefore new subtyping rules are needed to describe how subtyping relates with polymorphism:

$$\begin{array}{c}
\text{ALL} \\
\frac{\tau \triangleright \sigma}{\forall(\alpha) \tau \triangleright \forall(\alpha) \sigma} \\
\text{TRANS} \\
\frac{\tau \triangleright \sigma \quad \sigma \triangleright \rho}{\tau \triangleright \rho} \\
\text{INSTGEN} \\
\frac{\beta \notin \text{ftv}(\forall(\bar{\alpha}) \tau)}{\forall(\bar{\alpha}) \tau \triangleright \forall(\bar{\beta}) \tau[\bar{\alpha} \leftarrow \bar{\sigma}]} \\
\text{DISTRIB-R} \\
\frac{\alpha \notin \text{ftv}(\tau)}{\forall(\alpha) (\tau \rightarrow \sigma) \triangleright \tau \rightarrow \forall(\alpha) \sigma}
\end{array}$$

The congruence rule for polymorphic types ALL is unsurprising. Rule TRANS is uninteresting but it is needed in this presentation as it does not follow from other rules. The two interesting rules are INSTGEN and DISTRIB-R. Rule INSTGEN allows implicit toplevel type instantiation, as in ML: it means that a polymorphic type can be freely used at any of its instances. By contrast with ML, this rule can also be applied under arrow types at the appropriate variance. We write $\bar{\alpha}$ for a sequence of type variables (and similarly for types) with the understanding that $\forall(\bar{\alpha}) \tau$ stands for a sequence of quantifiers. Free type variables of τ are written $\text{ftv}(\tau)$. Notice, that although INSTGEN may generalize type variables that are introduced during type instantiation, it cannot generalize other type variables. Therefore, F_η still need a specific term typing rule for polymorphism introduction. Rule DISTRIB-R allows a quantifier to be pushed down on the right of arrow types when it does not appear on the left. The original rule allowed α to also occur in τ ; then the quantifier must be pushed down on both sides of the arrow, *i.e.* the right-hand side of the coercion becomes $(\forall(\alpha) \tau) \rightarrow \forall(\alpha) \sigma$. This more general rule is however derivable from DISTRIB-R and the other rules.

The Rule ARROW can be easily explained in F_η by the η -expansion of arrow types. Since F_η is closed by η -expansion, whenever a term a has some type $\tau \rightarrow \sigma$ in some context Γ , any type of its η -expansion $\lambda x. a x$ should also be a type of a . In particular, given $\tau' \triangleright \tau$ and $\sigma \triangleright \sigma'$, we may type the η -expansion by giving the

parameter x the type τ' and coerce the occurrence of x as an argument of a to the expected type τ which results in a value of type σ that can be coerced to σ' . Hence, the term a also has type $\tau' \rightarrow \sigma'$ in F_η .

What distinguishes type containment from traditional uses of subtyping is the inclusion in the subtyping relation of the implicit instantiation of quantifiers, which by congruence can be applied deeply inside terms. Unfortunately, this theoretical strength turns into a weakness for using F_η in practice since the type containment relation becomes undecidable.

Bounded quantification While F_η extends **Amber** with polymorphic types and can prove subtyping relations between polymorphic types, it does not allow to make subtyping assumptions about polymorphic type variables. Bounded quantification introduced by Cardelli and Longo (1991) in the language **Quest** and later in the language $F_{<}$ (Cardelli et al. 1994) solves this problem. Polymorphic type variables are introduced with an upper bound $\forall(\alpha <: \tau) \sigma$, whose instances are types of the form $\sigma[\alpha \leftarrow \rho]$ where ρ is a subtype of τ . The unbounded quantification of System F can be recovered as the special case where the bound is \top . Bounded polymorphism is quite expressive and has been used to model record subtyping and object oriented features.

The most interesting rule in $F_{<}$ is subtyping between bounded quantifications:

$$\frac{\text{FSUB} \quad \Gamma \vdash \tau' \triangleright \tau \quad \Gamma, \alpha \triangleright \tau' \vdash \sigma \triangleright \sigma'}{\Gamma \vdash \forall(\alpha <: \tau) \sigma \triangleright \forall(\alpha <: \tau') \sigma'}$$

Notice that subtyping holds only between two types having bounded quantification on both sides. That is, by contrast with F_η , type instantiation is not part of the subtyping relation and is made explicit at the level of terms. A good reason for this choice is to ease typechecking. Still, $F_{<}$ in its full generality is undecidable (Pierce 1994), but some restrictions of the rule, for instance where the bounds are identical on both sides, are decidable. Bounded polymorphism has also been extended to F-bounded polymorphism that allows the variable abstracted over to also appear in the bound, which is useful in object-oriented languages. The left-premise of Rule FSUB is then replaced by $\Gamma, \alpha \triangleright \tau' \vdash \alpha \triangleright \tau$.

Bounded polymorphism has been extensively used in many subsequent languages that combine polymorphism and subtyping, and in particular, to explore typechecking of objects (Abadi and Cardelli 1996).

Still, bounded polymorphism remains surprising in several ways. First, its formulation is asymmetric since type variables are introduced with an upper bound but no lower bound: dually, is there a use for lower bounds? Second, type variables have a unique bound: could the same variable have several bounds? Finally, while bounded polymorphism adds to F_η the ability to abstract over

subtyping, it does not yet generalize type containment since subtyping holds only for types with the same polymorphic structure; in particular, the implicit instantiation $\forall(\alpha <: \top) \tau \triangleright \tau[\alpha \leftarrow \sigma]$ that is allowed in F_η is not permitted in $F_{<}$.

Instance-bounded quantification Surprisingly, the absence of lower bounds in $F_{<}$ seemed to have not been a practical limitation and the question remained mostly theoretical for more than a decade. The need for lower bounds finally appeared for performing first-order type inference in the presence of second-order types. Predictable, efficient type inference is usually built on a notion of principal types, *i.e.* the ability to capture as a simple type (or type constraint) the set of all possible types of an expression. Lower bounds allow more expressions to have principal types.

Full type inference for System F is known for not having principal types because guessing types of function parameters often leads to an infinite set of solutions. However, if we restrict our ambition to not guess polymorphic types, but just propagate them, which is the goal of MLF (Le Botlan and Rémy 2009), there is still a problem to solve, namely the absence of principal solutions.

To illustrate this, consider the function `revapp` defined as $\lambda x. \lambda f. f x$ of type $\forall(\alpha) \forall(\beta) \alpha \rightarrow (\alpha \rightarrow \beta) \rightarrow \beta$ and its partial application to the identity function `id` of type $\forall(\gamma) \gamma^2$ where γ^2 stands for $\gamma \rightarrow \gamma$. What should be the type of `revapp id`? The type $\forall(\beta) ((\forall(\gamma) \gamma^2) \rightarrow \beta) \rightarrow \beta$, say τ_1 , obtained by keeping the identity polymorphic? Or the type $\forall(\gamma) \forall(\beta) (\gamma^2 \rightarrow \beta) \rightarrow \beta$, say τ_2 , obtained by instantiating `id` to γ^2 before the application and generalizing the result of the whole application afterwards? Unfortunately, no other type of the application `revapp id` is an instance of both of these two types in System F. (Nor in F_η ! but, conversely, both types have a common instance $\forall(\beta) (\forall(\gamma) (\gamma^2 \rightarrow \beta)) \rightarrow \beta$ in F_η .)

This dilemma is solved in MLF by introducing lower-bounded (or, rather, instance-bounded) polymorphism $\forall(\alpha :> \tau) \sigma$, which reads “*for all α that is an instance of τ , σ .*” The application `revapp id` can then be assigned the type $\forall(\alpha :> \forall(\gamma) \gamma^2) \forall(\beta) (\alpha \rightarrow \beta) \rightarrow \beta$, which happens to be principal in MLF. In particular, τ_1 can be obtained by inlining the bound which is permitted by instantiation in MLF; and τ_2 can be obtained by applying instantiation under the lower bound, generalizing γ afterwards which gives $\forall(\gamma) \forall(\alpha :> \gamma^2) \forall(\beta) (\alpha \rightarrow \beta) \rightarrow \beta$, and finally inlining the bound.

While in this regard MLF appears to be dual of $F_{<}$, it is still quite different: MLF does not use contravariance of arrow types as both $F_{<}$ and F_η do; conversely, MLF can freely instantiate polymorphic types, which $F_{<}$ cannot do: each of the three languages shares one feature with others but is still missing one of their key features.

$a, b ::= x \mid \lambda x. a \mid a a$	Terms
$\tau, \sigma ::= \alpha \mid \tau \rightarrow \tau \mid \perp \mid \top \mid \forall(\alpha : \kappa) \tau \mid \langle \rangle \mid \langle \tau, \tau \rangle \mid \pi_i \tau$	Types
$\kappa ::= \star \mid \{\alpha : \kappa \mid P\} \mid 1 \mid \kappa \times \kappa$	Kinds
$P ::= (\Delta \vdash \tau) \triangleright \tau \mid \exists \kappa \mid \top \mid P \wedge P$	Propositions
$\Gamma ::= \emptyset \mid \Gamma, (\alpha : \kappa) \mid \Gamma, (x : \tau)$	Environments

Figure 1: Syntax

Subtyping constraints The absence of multiple bounds in $F_{<}$ is somewhat surprising since ML extended with subtyping, say ML_{\leq} , naturally comes with (and requires the use of) subtyping constraints that mix arbitrary upper and lower bounds (Odersky et al. 1999). It uses constrained type schemes of the form $\forall(\alpha \mid C) \tau$ where C is a set of arbitrary but coherent subtyping constraints between simple types. Can we extend subtyping constraint to first-class polymorphic types?

2 The language F_{cc}

We have seen several type systems with different combinations of subtyping features, but none of them supersedes all others. In particular, F_{η} , $F_{<}$, MLF , and ML_{\leq} are pairwise incomparable. We now describe an extension of System F with coercion constraints, called F_{cc} , that combines all features together so that each of the languages above becomes a (still interesting) subset of F_{cc} . Here, we present a small subset of the language just to carry the main ideas underlying its design. We refer the reader to (Cretin and Rémy 2014b) for technical details.

Terms The language F_{cc} is implicitly typed for reasons that will be explained later—but this happens to be an advantage for conciseness: terms are just those of the untyped λ -calculus, whose notations are reminded in Fig. 1.

Although we focus on the syntactic presentation of F_{cc} hereafter, we assume that its semantics is given by full β -reduction. This is to build its type system on solid ground, without taking advantage of the evaluation strategy: it prevents from sweeping errors under λ 's just because their evaluation is delayed. Full β -reduction also models reduction of open terms. A practical language based on F_{cc} will eventually have a call-by-name or call-by-value evaluation strategy, which being a subset of full β -reduction, will remain sound.

Types, Kinds, and Propositions Types, defined on Figure 1, are simple types (type variables, arrow types, top, and bottom types) extended with con-

$$\begin{array}{c}
\text{TERMVAR} \\
\frac{(x : \tau) \in \Gamma}{\Gamma \vdash x : \tau}
\end{array}
\qquad
\begin{array}{c}
\text{TERMLAM} \\
\frac{\Gamma \vdash \tau : \star \quad \Gamma, (x : \tau) \vdash a : \sigma}{\Gamma \vdash \lambda x. a : \tau \rightarrow \sigma}
\end{array}
\qquad
\begin{array}{c}
\text{TERMAPP} \\
\frac{\Gamma \vdash a : \tau \rightarrow \sigma \quad \Gamma \vdash b : \tau}{\Gamma \vdash ab : \sigma}
\end{array}$$

$$\begin{array}{c}
\text{TERMCOER} \\
\frac{\Gamma, \Delta \vdash a : \tau \quad \Gamma \vdash (\Delta \vdash \tau) \triangleright \sigma}{\Gamma \vdash a : \sigma}
\end{array}$$

Figure 2: Term typing rules

strained polymorphic types of the form $\forall(\alpha : \kappa) \tau$ where κ is a kind that restricts the set of types in which α may range.

The kind \star is that of ordinary types, *e.g.* to form arrow types. The interesting kind is $\{\alpha : \kappa \mid P\}$ —the kind of types α of kind κ that satisfy the proposition P .

The grammars of types and kinds also include type tuples¹ classified by tuple kinds. Namely, type tuples are constructed from the empty tuple $\langle \rangle$ of kind 1 and pairs of types $\langle \tau_1, \tau_2 \rangle$ of kind $\kappa_1 \times \kappa_2$ when τ_i has kinds κ_i ; and type projection $\pi_i \tau$ to return the i 's component of kind κ_i of a tuple type of kind $\kappa_1 \times \kappa_2$. Tuple types are useful for capturing multiple binders but are not technically difficult: the main technical change is to consider types equal modulo the projection of tuples and closing equality by equivalence and congruence for all syntactical constructs. Thus, we will ignore tuples (which are grayed in Figure 1) in the rest of the technical overview.

Propositions are used to restrict sets of types. The most useful proposition is the coercion proposition $(\Delta \vdash \tau) \triangleright \sigma$ which states that there exists a coercion from type τ in some context extended with Δ , to type σ . For now, one may just consider the particular case of *simple* coercions where Δ is empty, in which case the proposition is abbreviated as $\tau \triangleright \sigma$ and just means that type τ can be coerced to type σ . The general case will be explained together with the term typing rules.

The proposition $\exists \kappa$ asserts that the kind κ is coherent (intuitively, that it is inhabited, but this is relative to its typing context). It is interesting that coherence can be internalized as a proposition. The proposition \top is the true proposition and the proposition $P_1 \wedge P_2$ is the conjunction of P_1 and P_2 .

Finally, typing environments Γ bind type variables to kinds and program variables to types. The letter Δ is used to range over environments that only bind type variables.

Term typing judgment ($\Gamma \vdash a : \tau$) Typing rules, given in Figure 2 are almost as simple as in Amber: they contain the typing rules of the simply typed

¹Not to be confused with the type of term tuples, which we do not include in this core version.

$$\begin{array}{c}
\text{TYPEPACK} \\
\frac{\Gamma, (\alpha : \kappa) \Vdash P \quad \Gamma \vdash \tau : \kappa \quad \Gamma \vdash P[\alpha \leftarrow \tau]}{\Gamma \vdash \tau : \{\alpha : \kappa \mid P\}}
\end{array}
\qquad
\begin{array}{c}
\text{TYPEUNPACK} \\
\frac{\Gamma \vdash \tau : \{\alpha : \kappa \mid P\}}{\Gamma \vdash \tau : \kappa}
\end{array}$$

Figure 3: Type judgment relation (excerpt)

λ -calculus plus the coercion typing rule, which differs from the **Amber Rule** in two important ways. First, the coercion judgment depends on the context Γ , as in $F_{<}$, so that coercion assumptions in the context Γ can be used to prove a coercion judgment such as $\Gamma \vdash \tau \triangleright \sigma$. More importantly, coercions are of the more general form $(\Delta \vdash \tau) \triangleright \sigma$ rather than just $\tau \triangleright \sigma$. As we can see in Rule **TERMCOER**, the context Δ contains additional type variable bindings that are added to the context Γ of the premise under which the coerced term must have type τ . The typical use of this generalization is in the judgment $\Gamma \vdash (\alpha : \kappa \vdash \tau) \triangleright \forall(\alpha : \kappa) \tau$, which holds whenever κ is coherent and $\forall(\alpha : \kappa) \tau$ is well-formed in Γ (see Rule **COERGEN** below). As a consequence, the rule for polymorphism introduction

$$\begin{array}{c}
\text{TERMGEN} \\
\frac{\Gamma \vdash \exists \kappa \quad \Gamma, (\alpha : \kappa) \vdash a : \tau}{\Gamma \vdash a : \forall(\alpha : \kappa) \tau}
\end{array}$$

is derivable by **TERMCOER** and **COERGEN**. It is remarkable that polymorphism introduction (as well as all erasable features of the type system) can be handled purely as a coercion typing rule and need no counterpart in term typing rules.

Well-formedness rules ($\Gamma \Vdash \kappa$ and $\Gamma \Vdash P$) The judgments $\Gamma \Vdash \kappa$ and $\Gamma \Vdash P$ state well-formedness of kinds and propositions. Besides syntactical checks, they are recursively scanning their subexpressions for all occurrences of coercion propositions $(\Delta \vdash \tau) \triangleright \sigma$ to ensure that Δ , τ , and σ are well-typed, as described by the following rule:

$$\frac{\Gamma \vdash \Delta \quad \Gamma, \Delta \vdash \tau : \star \quad \Gamma \vdash \sigma : \star}{\Gamma \Vdash (\Delta \vdash \tau) \triangleright \sigma}$$

In particular, the auxiliary judgment $\Gamma \vdash \Delta$, which treats Δ as a telescope, means that each binding $\alpha : \kappa$ of Δ is well-formed in the typing context that precedes it, which also implies that the kind κ is coherent in any such binding relatively to its typing context.

Kinding judgment ($\Gamma \vdash \tau : \kappa$) An excerpt of kinding rules is given in Figure 3, but most rules have been omitted. Rule **TYPEUNPACK** states that whenever a type τ is known to be of a constrained kind $\{\alpha : \kappa \mid P\}$, it is also of the kind κ , indeed.

$$\frac{\text{PROPRES}}{\Gamma \vdash \tau : \{\alpha : \kappa \mid P\}} \quad \frac{\text{PROPEXI}}{\Gamma \vdash \tau : \kappa} \quad \frac{}{\Gamma \vdash \exists \kappa}$$

Figure 4: Proposition judgment relation (excerpt)

$$\begin{array}{c} \text{COERREFL} \\ \frac{}{\Gamma \vdash \tau \triangleright \tau} \\ \text{COERTOP} \\ \frac{}{\Gamma \vdash \tau \triangleright \top} \\ \text{COERBOT} \\ \frac{}{\Gamma \vdash \perp \triangleright \tau} \\ \text{COERTRANS} \\ \frac{\Gamma, \Delta' \vdash (\Delta \vdash \tau) \triangleright \tau' \quad \Gamma \vdash (\Delta' \vdash \tau') \triangleright \tau''}{\Gamma \vdash (\Delta', \Delta \vdash \tau) \triangleright \tau''} \\ \text{COERWEAK} \\ \frac{\Gamma \vdash (\Delta \vdash \tau) \triangleright \sigma}{\Gamma \vdash \tau \triangleright \sigma} \\ \text{COERARR} \\ \frac{\Gamma \vdash \tau : \star \quad \Gamma, \Delta \vdash \tau \triangleright \tau' \quad \Gamma \vdash (\Delta \vdash \sigma') \triangleright \sigma}{\Gamma \vdash (\Delta \vdash \tau' \rightarrow \sigma') \triangleright \tau \rightarrow \sigma} \\ \text{COERGEN} \\ \frac{\Gamma \vdash \exists \kappa \quad \Gamma, (\alpha : \kappa) \vdash \tau : \star}{\Gamma \vdash (\alpha : \kappa \vdash \tau) \triangleright \forall (\alpha : \kappa) \tau} \\ \text{COERINST} \\ \frac{\Gamma, (\alpha : \kappa) \vdash \tau : \star \quad \Gamma \vdash \sigma : \kappa}{\Gamma \vdash \forall (\alpha : \kappa) \tau \triangleright \tau[\alpha \leftarrow \sigma]} \end{array}$$

Figure 5: Coercion judgment relation

TYPEPACK shows that, conversely, knowing that τ has kind κ , one must still prove that $P[\alpha \leftarrow \tau]$ is satisfied to be able to consider that τ has the constrained kind $\{\alpha : \kappa \mid P\}$.

Proposition judgment ($\Gamma \vdash P$) The two most interesting rules for propositions are given in Figure 4. Rule PROPRES means that a type of a constrained kind $\{\alpha : \kappa \mid P\}$ satisfies the proposition P (where τ has been substituted for α). Rule PROPEXI means that one must exhibit a type τ of kind κ to ensure that the kind κ is coherent in its typing context Γ .

Coercion propositions ($\Gamma \vdash (\Delta \vdash \tau) \triangleright \sigma$) We have separated the rules for coercions in Figure 5 although coercions are just a particular case of propositions. Some rules are slightly obfuscated by the coercion typing context Δ that has to be added in some premises. These rules can first be read in the particular case where Δ is the empty context, so we have grayed Δ to help with this reading.

Rule COERREFL, COERTOP, and COERBOT are obvious and can be skipped. Rule COERTRANS is also standard—the Δ 's just need to be appropriately concatenated.

Rule COERWEAK implements a form of weakening: it tells that if any term of typing² $\Gamma, \Delta \vdash \tau$ has typing $\Gamma \vdash \sigma$, then any term of typing $\Gamma \vdash \tau$ also has typing $\Gamma \vdash \sigma$. Weakening is required as it would not be derivable from the other rules if

²We say that a term a has typing $\Gamma \vdash \tau$ if the typing judgment $\Gamma \vdash a : \tau$ holds.

we removed it from the definition.

Rule COERARR is the usual contravariant rule for arrows when Δ is empty. Otherwise, the rule can be understood by looking at the η -expansion context $\lambda x. (\square x)$ for the arrow type. Placing a term with typing $\Gamma, \Delta \vdash \tau' \rightarrow \sigma'$ in the hole, we may give $\lambda x. (\square x)$ the typing $\Gamma \vdash \tau \rightarrow \sigma$ provided a coercion of type $\Gamma, \Delta \vdash \tau \triangleright \tau'$ is applied around x . The result of the application has typing $\Gamma, \Delta \vdash \sigma'$ which can in turn be coerced to $\Gamma \vdash \sigma$ if there exists a coercion of type $\Gamma \vdash (\Delta \vdash \sigma') \triangleright \sigma$. Thus, the η -expansion has typing $\Gamma \vdash \tau \rightarrow \sigma$.

Notice that COERARR is the only η -expansion rule, since the arrow is the only computational type constructor in our core language. For each computational type constructor that we would add to the language, we would need a corresponding η -expansion coercion rule. Erasable type constructors need not have an η -expansion coercion rule, since these are derivable as their introduction and elimination rules are already coercions.

Rule COERGEN implements type generalization, but as a coercion rule. It says that $\Gamma \vdash (\alpha : \kappa \vdash \tau) \triangleright \forall(\alpha : \kappa) \tau$ is a valid coercion as long as kind κ is coherent in Γ and type τ has kind \star in $\Gamma, (\alpha : \kappa)$. This coercion makes TERMGGEN derivable as explained above.

Similarly, COERINST is the counterpart of type instantiation coercion in F_η : it says that $\Gamma \vdash \forall(\alpha : \kappa) \tau \triangleright \tau[\alpha \leftarrow \sigma]$ is a valid coercion as long as τ has kind \star in $\Gamma, (\alpha : \kappa)$ and σ has kind κ in Γ . Since COERGEN is a coercion rule, we do not need the more involved version of F_η that performs generalization afterwards.

Notice that there is no counterpart to DISTRIB-R since it is derivable by a combination of type generalization, type instantiation, and η -expansion.

Adding recursive types and coinduction

The full language F_{cc} also has recursive types. These are indeed present in the language Amber, $F_{<}$, *etc.* as they are unavoidable in a programming language.

We thus extend the grammar of types with the production $\mu\alpha \tau$, with the usual restriction on well-foundedness of recursive types that we will not detail here. We write $\alpha \mapsto \tau : \mathbf{wf}$ to mean that the function $\alpha \mapsto \tau$ is well-founded and can be used to form the recursive type $\mu\alpha \tau$.

Since types are implicit, we chose equi-recursive types, *i.e.* the equality of recursive types is witnessed by coercions, which are implicit. We add two coercions to witness folding and unfolding of recursive types:

$$\frac{\text{COERUNFOLD} \quad \alpha \mapsto \tau : \mathbf{wf} \quad \Gamma, (\alpha : \star) \vdash \tau : \star}{\Gamma \vdash \mu\alpha \tau \triangleright \tau[\alpha \leftarrow \mu\alpha \tau]} \quad \frac{\text{COERFOLD} \quad \alpha \mapsto \tau : \mathbf{wf} \quad \Gamma, (\alpha : \star) \vdash \tau : \star}{\Gamma \vdash \tau[\alpha \leftarrow \mu\alpha \tau] \triangleright \mu\alpha \tau}$$

To reason with recursive types we also add coinduction in propositions. In

order to do so, we change the judgment $\Gamma \vdash P$ to $\Gamma; \Theta \vdash P$ where Θ is used to collect coinductive hypotheses. We introduce two new rules `PROPFIX` and `PROPVAR` to allow reasoning by coinduction.

$$\frac{\text{PROPFIX} \quad \Gamma \Vdash P \quad \Gamma; \Theta, P \vdash P}{\Gamma; \Theta \vdash P} \quad \frac{\text{PROPVAR} \quad P^\dagger \in \Theta}{\Gamma; \Theta \vdash P} \quad \frac{\text{COERARR} \quad \Gamma, \Delta; \Theta^\dagger \vdash \tau \triangleright \tau' \quad \Gamma; \Theta^\dagger \vdash (\Delta \vdash \sigma') \triangleright \sigma}{\Gamma; \Theta \vdash (\Delta \vdash \tau' \rightarrow \sigma') \triangleright \tau \rightarrow \sigma}$$

Rule `PROPFIX` allows the judgment $\Gamma; \Theta \vdash P$ to be proved under the additional coinductive hypothesis P (provided P is well-formed in Γ). Of course, coinductive hypotheses must be guarded before they can be used. This is implemented by tagging guarded propositions in Θ that are then ready for coinductive use with \dagger . Hence, rule `PROPVAR` can prove P only if P appears guarded, *i.e.* as P^\dagger , in the coinductive context Θ .

Premises of the rule `COERARR` use only subterms of the arrow type, so the rule acts as a guard for all coinductive assumptions, therefore all propositions of Θ are tagged in its premises. This is only the case of expansion rules for computational type constructors, which have a counterpart in terms. Currently, `COERARR` is the only such rule. All other rules just transport Θ unchanged from the premise to the conclusion.

Reasoning by induction makes the following standard rules for equi-recursive types derivable (we write $\tau \triangleleft \sigma$ for a pair of simple coercions $\tau \triangleright \sigma$ and $\sigma \triangleright \tau$):

$$\frac{\text{COERPERIOD} \quad \alpha \mapsto \sigma : \mathbf{wf} \quad \Gamma; \Theta \vdash (\tau_i \triangleleft \sigma [\alpha \leftarrow \tau_i])^{i \in \{1,2\}}}{\Gamma; \Theta \vdash \tau_1 \triangleright \tau_2} \quad \frac{\text{COERETAMU} \quad \Gamma, (\alpha, \beta, \alpha \triangleright \beta); \Theta \vdash \tau \triangleright \sigma}{\Gamma; \Theta \vdash \mu \alpha \tau \triangleright \mu \beta \sigma}$$

Interestingly, the proof for `COERPERIOD` requires reinforcement of the coinduction hypothesis since we need $\tau_1 \triangleleft \tau_2$ and not just $\tau_1 \triangleright \tau_2$ in the coinduction hypothesis.

3 Strength and weaknesses of F_{cc}

Soundness The type system of F_{cc} is sound for the full β -reduction semantics. Type soundness is not proved syntactically for reasons explained next, but semantically, by interpreting types as sets of terms. As a consequence of the presence of general recursive types (*i.e.* we do not restrict to positive recursion), we use a step-indexed technique. Unfortunately, the usual technique of Appel and McAllester (2001) does not apply to a full β -reduction setting. We propose a new technique where indexes that are traditionally outside terms are placed directly on terms and are transformed during reduction. See (Cretin and Rémy 2014a) for details.

Termination As a sanity check, the reduction of well-typed programs always terminates in the absence of recursive types and coinduction.

(Lack of a proof of) Subject reduction The reason not to do a syntactic proof is that we do not know how to prove subject reduction for F_{cc} . The problem is that the type system is either too expressive or too weak: it allows to type programs that are indeed safe, but with involved non local coercion constraints that cannot be easily traced during reduction.

Doing a syntactic proof would amount to having an explicitly typed version of the language and reduction rules for explicitly typed terms that preserve well-typedness; moreover, reduction of explicitly typed terms must be in bisimulation with the reduction of implicitly typed ones. The main obstacle is that, in the general case, abstract coercions may appear in the middle of a redex. Explaining why this is difficult in the general case is a bit tricky, as solving one issue immediately raises another one—see Cretin and Rémy (2012) for details. Quite interestingly, this configuration can never occur when we restrict to abstract coercions that are parametric in either their domain or their codomain, *i.e.* coercions of the form $\alpha \triangleright \tau$ or $\tau \triangleright \alpha$. Remarkably, these two subcases coincide with bounded quantification and instance-bounded quantification. Under this restriction, we can design an explicitly typed language that enjoys subject reduction (Cretin and Rémy 2012).

(Lack of a good) Surface language Since the type system is implicit, it is undecidable, of course. What is a good surface language for F_{cc} is still an open question. It is always possible to be fully explicit by introducing term syntax for describing typing derivations in source terms, hence turning type inference into an easy checking process. However, this would not only contain type annotations but also full coercion bodies (information which is typically inferred in languages such as $F_{<}$) and of coercion coherence proofs (information which is usually obtained by construction). Programming at this level of detail would be too cumbersome for the programmer. Notice however, that F_η already suffers from a similar problem, since its coercion relation is undecidable.

This issue can be addressed in two directions. Remaining within F_{cc} , we may apply partial type inference techniques and hope that sufficient type and coercion information can be reconstructed. It would also be interesting to look for subsystems of F_{cc} that compromise expressiveness for a smaller amount of annotations. The restriction to coercions that are parametric in either their domain or their codomain is one such solution. Are there other sweet spots?

Expressiveness As announced earlier, F_{cc} contains F_η , $F_{<}$, MLF , and ML_{\leq} as sublanguages. This confirms that all their features are compatible and can safely be combined together.

Writing $\forall(\alpha \mid P)\tau$ for $\forall(\alpha : \{\alpha : \star \mid P\})\tau$, bounded quantification and instance-bounded quantification can be encoded in F_{cc} as $\forall(\alpha \mid \alpha \triangleright \tau)\sigma$ and $\forall(\alpha \mid \tau \triangleright \alpha)\sigma$. It is then not difficult to see that the typing rules of $F_{<}$ (including F-Bounded quantification) and of MLF are derivable.

The notation is generalized to bindings, writing $(\bar{\alpha} \mid P)$ for $(\alpha : \{\alpha : \star^n \mid P\})$ where n is the size of $\bar{\alpha}$ and, by abuse of notation, let us write α_i for $\pi_i \alpha$ when α_i is the i ' component of a sequence $\bar{\alpha}$. Then, a constrained typing judgment $\Gamma \vdash e : \tau \mid C$ in ML_{\leq} can be seen as the F_{cc} judgment $(\bar{\alpha} \mid C), \Gamma \vdash e : \tau$ where $\bar{\alpha}$ are free variables of Γ , C , and τ .

We have only presented a core subset of F_{cc} . The full language (Cretin and Rémy 2014a) also contains products. Sum types could be also be easily added. Existential types can be emulated by their CPS encoding. The language of propositions contains polymorphic propositions $\forall(\alpha : \kappa)P$, and could also be enriched with other propositions.

Incoherent coercions In the subset of F_{cc} described above, coercions are always coherent relative to the typing context in which they are used. More precisely, when an expression a has type $\forall(\alpha : \kappa)\tau$ in Γ , it is always the case that $\Gamma \vdash \exists \kappa$. This is necessary because type abstraction does not have any counterpart in terms and, in particular, does not block the evaluation of a which may proceed immediately. Without coherence, one could abstract over the absurd kind constraint $\top \triangleright \perp$ and be able to type any program.

However, there are situations in which abstraction over incoherent coercions would be useful. First, the coercion may be coherent only for *some* instances of the typing context. This is typically the case in the presence of GADTs. Second, coherence at the abstraction point is often harder to prove than at instantiation points where types have been specialized. For these reasons, we have also extended F_{cc} with incoherent coercion abstractions. Of course, this abstraction must now block the evaluation and therefore have a counterpart in terms. Interestingly, this allows to model GADTs as incoherent coercions. Incoherent abstractions are indeed used in FC (Weirich et al. 2011), the intermediate language of Haskell. See (Cretin and Rémy 2014a) for further details.

Conclusions

We have given a tour of coercion constraints and shown how they can be used to explain several type systems that had been designed separated for different

purposes, but all around some variations on the notion of subtyping. There are an amazingly large number of interesting works on subtyping, many of which actually initiated by Cardelli, and we could unfortunately just select a few citations among all the relevant ones. As for coercions, the idea is not at all new. There are in fact several notions of coercions and many works on type systems have already used coercions as a tool or studied them on their own. So, many more references could have been included here as well. We refer the reader to (Cretin and Rémy 2014a) for a more thorough treatment of related works.

Coercion constraints can factor many type features of programming languages. This allows sharing an important part of their meta-theoretical studies, such as type soundness. It also opens the door to new combinations of features. Besides, it helps separate the computational and erasable parts of type systems. We thus believe that they are an interesting framework for designing and studying type systems.

However, F_{cc} still lacks a good surface language for the programmer as well an explicitly-typed calculus to be used as an internal language in a compiler. While partial type inference techniques could be used for the surface language, finding an explicit version of coercion constraints that enjoys subject reduction without sacrificing expressiveness seems much more challenging.

References

- M. Abadi and L. Cardelli. *A Theory of Objects*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 1st edition, 1996. ISBN 0387947752.
- A. W. Appel and D. McAllester. An indexed model of recursive types for foundational proof-carrying code. *ACM Transactions on Programming Languages and Systems.*, 23(5), Sept. 2001.
- L. Cardelli. A semantics of multiple inheritance. In *Proc. Of the International Symposium on Semantics of Data Types*, pages 51–67, New York, NY, USA, 1984. Springer-Verlag New York, Inc. ISBN 3-540-13346-1. URL <http://dl.acm.org/citation.cfm?id=1096.1098>.
- L. Cardelli and G. Longo. A semantic basis for quest. *J. Funct. Program.*, 1(4): 417–458, 1991.
- L. Cardelli, S. Martini, J. C. Mitchell, and A. Scedrov. An extension of system f with subtyping. *Information and Computation*, 109(1/2):4–56, 1994.
- J. Cretin and D. Rémy. On the power of coercion abstraction. In *Proceedings of the annual symposium on Principles Of Programming Languages*, 2012.

- J. Cretin and D. Rémy. System F with Coercion Constraints. Rapport de recherche RR-8456, INRIA, Jan. 2014a. URL <http://hal.inria.fr/hal-00934408>.
- J. Cretin and D. Rémy. System F with Coercion Constraints. In *Logic In Computer Science (LICS)*, July 2014b. To appear.
- D. Le Botlan and D. Rémy. Recasting MLF. *Information and Computation*, 207(6), 2009.
- J. C. Mitchell. Type inference and type containment. In *Proc. Of the International Symposium on Semantics of Data Types*, pages 257–277, New York, NY, USA, 1984. Springer-Verlag New York, Inc. ISBN 3-540-13346-1. URL <http://dl.acm.org/citation.cfm?id=1096.1106>.
- J. C. Mitchell. Polymorphic type inference and containment. *Information and Computation*, 2/3(76), 1988.
- M. Odersky, M. Sulzmann, and M. Wehr. Type inference with constrained types. *Theory and Practice of Object Systems*, 5(1):35–55, 1999.
- B. C. Pierce. Bounded quantification is undecidable. *Information and Computation*, 112(1):131–165, July 1994.
- J. C. Reynolds. Using category theory to design implicit conversions and generic operators. In N. D. Jones, editor, *Semantics-Directed Compiler Generation*, volume 94 of *Lecture Notes in Computer Science*, pages 211–258, Berlin, 1980. Springer-Verlag.
- S. Weirich, D. Vytiniotis, S. Peyton Jones, and S. Zdancewic. Generative type abstraction and type-level computation. In *Proceedings of the 38th annual ACM SIGPLAN-SIGACT symposium on Principles of programming languages, POPL'11*, 2011.