

UNIVERSITÉ PARIS DIDEROT (PARIS 7)
ÉCOLE DOCTORALE 386 : SCIENCES MATHÉMATIQUES DE PARIS CENTRE
ÉQUIPE GALLIUM, INRIA

DOCTORAT
INFORMATIQUE

Coercions effaçables :
une approche unifiée des systèmes de types

english title:

Erasable coercions:
a unified approach to type systems

Julien CRETIN

(final version)

Thèse dirigée par Didier RÉMY
et soutenue le 30 janvier 2014 devant le Jury composé de :

Président	Roberto DI COSMO
Rapporteurs	Mariangiola DEZANI-CIANCAGLINI Stephanie WEIRICH
Examineurs	Jean-Christophe FILLIÂTRE Paul-André MELLIÈS
Directeur	Didier RÉMY

Abstract

Functional programming languages, like OCaml or Haskell, rely on the λ -calculus for their core language. Although they have different reduction strategies and type system features, their proof of soundness and normalization (in the absence of recursion) should be factorizable. This thesis does such a factorization for theoretical type systems featuring recursive types, subtyping, bounded polymorphism, and constraint polymorphism. Interestingly, soundness and normalization for strong reduction imply soundness and normalization for all usual strategies. Our observation is that a generalization of existing coercions permits to describe all type system features stated above in an erasable and composable way. We illustrate this by proposing two concrete type systems: first, an explicit type system with a restricted form of coercion abstraction to express subtyping and bounded polymorphism; and an implicit type system with unrestricted coercion abstraction that generalizes the explicit type system with recursive types and constraint polymorphism—but without the subject reduction property. A side technical result is an adaptation of the step-indexed proof technique for type-soundness to calculi equipped with a strong notion of reduction.

Remerciements

Avant de commencer, j'aimerais remercier les personnes sans qui je n'aurais pu terminer cette thèse. Il y a tout d'abord les personnes qui m'ont initié à l'informatique, en particulier Gilles Dowek, Luc Maranget, François Pottier et Benjamin Werner lors de ma formation à l'X. Je remercie également Benjamin Pierce, Nate Foster et Michael Greenberg avec qui j'ai découvert l'environnement de la recherche pendant mon stage à UPenn. C'est d'ailleurs pendant ce stage, grâce à Davi Barbosa et suite à une discussion avec Mike, que j'ai découvert Haskell qui finira par devenir un motif récurrent dans ma thèse.

Je remercie aussi mes professeurs du MPRI qui m'ont beaucoup apporté. C'est d'ailleurs pendant le MPRI que j'ai rencontré Didier qui a successivement été mon professeur, mon directeur de stage de master, et enfin mon directeur de thèse. Je le remercie tout particulièrement pour son suivi et ses idées.

Je remercie Inria, Gallium et Xavier Leroy pour mon environnement de travail à Rocquencourt. Les navettes ont fait en sorte que ce ne soit pas si loin de Paris. Et la verdure et l'espace en ont fait un cadre de travail très agréable. Je remercie en particulier Nicolas Pouillard qui a été mon voisin de bureau pendant ma première année de thèse et avec qui j'ai beaucoup appris par nos discussions sur Haskell, Agda, les lieux, Bitcoin, etc.

Je remercie Simon Peyton Jones, Stephanie Weirich et Dimitrios Vytiniotis pour m'avoir offert une pause dans toute cette théorie. J'ai pu en particulier découvrir en détail le fonctionnement d'un compilateur pour Haskell.

J'ai passé une grande partie de ma dernière année de thèse à écrire mes preuves en Coq, un outil que je ne maîtrise pas. L'aide de Jacques-Henri Jourdan, Thomas Braibant et François Pottier m'a été très utile et je les en remercie grandement. Indirectement, les idées d'Arthur Charguéraud sur la formalisation en Coq m'ont aussi été utiles.

Je remercie mes deux rapporteurs Mariangiola Dezani-Ciancaglini et Stephanie Weirich pour leur lecture attentive et leurs commentaires enrichissants. Je remercie aussi Gabriel Scherer pour sa lecture anticipée et ses nombreux commentaires qui m'ont permis d'éditer la première version du manuscrit. Je remercie enfin Thibaut Balabonski pour son aide précieuse à la mise au point de mes slides¹.

Il y a enfin toutes les personnes sans lien direct avec la science de cette thèse mais qui ont tout de même contribué à son déroulement. Je remercie mes parents, ma sœur, ma copine, ma troupe et mes amis pour les bons moments même s'ils ont parfois pu déborder sur ma thèse :-). Merci à Maëlle, Marie, Tomô, Barbara, Pauline, Muriel, Sophie, Stéphanie, Louis, Manu, Jacques, Antoine (RIP), Xavier, Flora, Claire, Xavier, Amandine, Mica, Guillaume, Chloé, Séb, Anakin, et ceux que j'oublie.

¹disponibles sur mon site de thèse : <http://phd.ia0.fr/slides.pdf>

Résumé

Les *systèmes de types* permettent de vérifier qu'un programme vérifie certaines propriétés avant toute exécution. La propriété de programme la plus commune est la correction : un programme est correct si son exécution ne rencontre pas d'erreurs. Une autre propriété intéressante est la terminaison, car les algorithmes sont des programmes corrects qui terminent : ils renvoient un résultat après un temps fini. Les systèmes de types classifient les programmes avec des invariants de comportement pour rejeter ceux qui sont potentiellement incorrects et conserver exclusivement ceux qui sont corrects. Les systèmes de types sont rarement complets : ils rejettent parfois des programmes corrects. Mais ils sont toujours corrects : ils ne valident que des programmes corrects. On dit qu'ils sont conservatifs. Puisque les invariants des systèmes de types précisent le comportement d'un programme, ils peuvent aussi servir de documentation.

Puisqu'un système de types classe les programmes, il repose donc sur un langage de programmation. En revanche, un langage de programmation n'a pas besoin d'être muni d'un système de types. Et lorsqu'il est muni d'un système de types, celui-ci n'est pas exclusif. Il arrive d'ailleurs qu'un même langage de programmation ait plusieurs systèmes de types. C'est par exemple le cas du λ -calcul (Chapitre 2). Le λ -calcul simplement typé (Section 3.1), System F (Section 3.2), System F_η (Section 3.4), MLF (Section 3.5), System $F_{<}$ (Section 3.6), pour n'en citer que quelques uns, sont des systèmes de types du λ -calcul. Chacun de ces systèmes est indépendamment prouvé correct et normalisant (tous les programmes terminent). Il serait donc intéressant d'étudier si ces systèmes de types ne sont pas unifiaibles afin de factoriser leur preuve de correction et de terminaison en l'absence de récursion.

Comme les systèmes de types sont conservatifs, il existe une relation d'ordre entre les systèmes d'un même langage de programmation. Un système de types T contient un autre système S si tous les programmes acceptés par S le sont aussi par T . Les propriétés de T peuvent donc être récupérées par S . Si un programme est accepté par S , il est accepté par T , et vérifie donc les propriétés de T (par exemple la correction ou la terminaison). Dans les systèmes du Chapitre 3, le système F_η contient le système F, qui contient à son tour le λ -calcul simplement typé. Il suffit donc de prouver les propriétés qui nous intéressent pour le système F_η .

Cette thèse présente un framework de coercions qui unifie les systèmes du Chapitre 3, et qui les contient donc. Les coercions peuvent être vues comme une extension de la notion de sous-typage des types aux typings effaçables (une paire d'un environnement effaçables et d'un type) comme décrit dans la Section Sous-typage. Ce travail distingue les types de calcul des types effaçables. Les premiers ont à faire au calcul et sont donc introduits et éliminés dans le jugement de typage des termes. Chaque règle correspond à un nœud du calcul. En revanche, les types effaçables ont à voir avec le typage. Ils sont introduits et éliminés avec des coercions et donc de manière effaçable car les coercions sont effaçables (Section Bisimulation). Cette distinction se voit clairement dans les règles de typage des termes. Les figures 4.10 et 5.10

contiennent six règles de calcul pour les types de calcul et une règle de retypage `TERMCOER` pour tous les types effaçables.

Ce document étudie en particulier l'abstraction de coercion (Section Abstraction de coercions). Celle-ci est vue comme une abstraction sur du contenu effaçable, à savoir les types et les coercions. Cette abstraction doit être cohérente pour être effaçable. Certaines abstraction sont naturellement cohérentes, comme le polymorphisme de Système F ou le polymorphisme borné de MLF et Système $F_{<}$, alors que d'autres abstractions requièrent une preuve de cohérence, comme le polymorphisme contraint de ML avec contraintes. Nous présentons deux systèmes de types : Système F_t^p (Chapitre 4) où les abstractions sont naturellement cohérentes et Système F_{cc} (Chapitre 5) où les abstractions requièrent une preuve de cohérence.

Pour conclure, cette thèse présente un framework de coercions effaçables pour décrire les caractéristiques effaçables des systèmes de types de manière unifiée. Les types de calcul sont introduits et éliminés par des termes alors que les types effaçables le sont par des coercions effaçables. Les types de calcul étudiés sont : la flèche et le produit pour observer la correction, et le polymorphisme incohérent pour les GADTs. Les types effaçables étudiés sont : les extrêmes top et bottom pour fermer la relation de coercion, le polymorphisme cohérent sous ses formes non contrainte, bornée et contrainte, et les types récursifs. Ces caractéristiques dans un framework avec eta-expansion et coinduction, permettent d'inclure Système F_η , MLF, Système $F_{<}$ et ML avec contraintes. Cela permet par exemple d'utiliser côte à côte l'inférence de MLF et celle de ML avec contraintes en utilisant les types de ML pour l'interface.

Bisimulation

Un système de types est dit explicite, lorsqu'il est muni d'une notion d'objets explicites pour que les dérivations soient dirigées par la syntaxe. Par exemple, les termes explicites contiennent les annotations de types et de coercions nécessaires pour reconstruire leur dérivation de typage à partir de leur environnement initial.

Lorsque les termes explicites sont munis d'une notion de réduction explicite, une question naturelle est d'étudier le lien entre cette réduction et la réduction implicite. Ce lien est fait à travers la fonction d'effacement des annotations et permet de montrer les lemmes de préservation du typage et de progrès en utilisant leurs analogues explicites. Les versions explicites sont plus simples à prouver, car la réduction explicite décrit la preuve de préservation du typage. Ce lien correspond à une bisimulation entre les deux réductions et correspond à la notion d'effaçabilité. Les annotations de types et de coercions sont effaçables en ce sens.

L'énoncé de ce lemme requiert de distinguer les pas de calcul des pas effaçables dans la réduction explicite. La réduction implicite simule les pas de calcul de la réduction explicite et ignore les pas effaçables. Et réciproquement, la réduction explicite simule la réduction implicite : un pas de réduction implicite correspond à un nombre fini de pas effaçables suivi d'un pas de calcul. Voir la Section 4.4.4 pour plus de détails.

Sous-typage

Le sous-typage permet de voir un terme de type τ avec le type σ sous un environnement Γ pourvu qu'il existe une dérivation de $\Gamma \vdash \tau \triangleright \sigma$, à savoir que τ soit un sous-type de σ sous Γ . Cette relation est une relation d'ordre et elle est donc réflexive et transitive. Pour illustrer simplement le sous-typage, il suffit d'enrichir les types avec deux extrêmes, top \top et bottom \perp .

Tous les types sont plus petit que top, à savoir $\Gamma \vdash \tau \triangleright \top$, et tous les types sont plus grands que bottom, à savoir $\Gamma \vdash \perp \triangleright \tau$. On peut étendre la relation de sous-typage par congruence. En particulier, si τ est un sous-type de τ' sous Γ et σ' est un sous-type de σ sous Γ , alors $\tau' \rightarrow \sigma'$ est un sous-type de $\tau \rightarrow \sigma$ sous Γ (voir Section Eta-expansion pour plus de détails). L'inversion de sens pour le type de l'argument vient de la contravariance de la flèche sur son argument.

Les coercions effaçables peuvent être vues comme une extension du sous-typage des types aux typings effaçables (paires d'un environnement effaçable et d'un type), ou bien comme une extension du sous-typage qui s'autorise à étendre l'environnement avec des lieux effaçable. Les coercions effaçables permettent de voir un terme de type τ sous l'environnement étendu Γ, Σ , avec le type σ sous l'environnement Γ , pourvu qu'il existe une dérivation de $\Gamma \vdash (\Sigma \vdash \tau) \triangleright \sigma$, c'est-à-dire que le typing effaçable $\Sigma \vdash \tau$ soit inclus dans le type σ sous l'environnement Γ . L'extension d'environnement Σ ne peut contenir que des lieux effaçables (de types ou de coercions). On appelle ces environnements des environnements effaçables. Il est évident que la relation de sous-typage $\Gamma \vdash \tau \triangleright \sigma$ est incluse dans la relation de coercion $\Gamma \vdash (\Sigma \vdash \tau) \triangleright \sigma$ en prenant l'environnement vide \emptyset pour l'environnement effaçable Σ de la coercion.

L'illustration la plus simple des coercions effaçables est le polymorphisme. Tout comme les règles d'introduction et d'élimination de top et bottom étaient exprimées dans la relation de sous-typage, les règles d'introduction et d'élimination du polymorphisme sont exprimables dans la relation de coercion. La règle d'élimination était déjà une règle de sous-typage (et donc de coercion) : le typing effaçable $\emptyset \vdash \forall \alpha \tau$ est coercible vers sesinstanciations $\tau[\alpha/\sigma]$ sous l'environnement Γ pour tout type σ bien formé sous Γ . La règle d'introduction du polymorphisme, quant à elle, ne peut être exprimée que dans la relation de coercion : le typing effaçable $\emptyset, \alpha \vdash \tau$ est coercible vers $\forall \alpha \tau$ sous Γ .

$$\begin{array}{c} \text{ELIM} \\ \frac{\Gamma \vdash \sigma \text{ type}}{\Gamma \vdash (\emptyset \vdash \forall \alpha \tau) \triangleright \tau[\alpha/\sigma]} \end{array} \qquad \begin{array}{c} \text{INTRO} \\ \Gamma \vdash (\emptyset, \alpha \vdash \tau) \triangleright \forall \alpha \tau \end{array}$$

Eta-expansion

Une caractéristique très importante du sous-typage et donc des coercions est de pouvoir appliquer une règle en profondeur dans un type en respectant la variance. Ce mécanisme est possible par les règles de congruence. Ce sont elles qui définissent la variance des constructeurs de type. Ces règles de congruence découlent de la notion d'eta-expansion et donc des règles d'introduction et d'élimination des constructeurs de types.

Il faut noter que les règles de congruence ne sont nécessaires que pour les types de calcul, car elles sont dérivables pour les types effaçables. En effet, puisque les règles d'introduction et d'élimination des types effaçables sont dans la relation de coercion, il suffit de procéder à une eta-expansion effaçable. Par exemple, dans le cas du polymorphisme, on peut montrer que le typing effaçable $\Sigma \vdash \forall \alpha \tau$ est coercible vers $\forall \alpha \sigma$ sous Γ pourvu que Σ ne mentionne pas α et que le typing $\Sigma \vdash \tau$ soit coercible vers σ sous l'environnement étendu Γ, α . Il suffit pour cela de composer une élimination, la coercion interne, et une introduction.

Les règles de congruence des types de calcul reposent sur une eta-expansion non effaçable et doivent donc être rajoutées dans la relation de coercions. Prenons le cas de la congruence de la flèche, dont l'eta-expansion s'écrit $\lambda x \square x$. Il est possible de placer des coercions à deux endroits, sur la variable et sur le corps de l'abstraction, qui correspondent respectivement à la

coercion sur le type de l'argument et le type de retour de la flèche. On constate alors que la coercion de la variable se trouve sous la portée des variables effaçables liées par la coercion du corps de l'abstraction. Si cette dernière coercion s'écrit $\Gamma \vdash (\Sigma \vdash \sigma') \triangleright \sigma$, alors la coercion de la variable $\tau \triangleright \tau'$ opère sous l'environnement étendu Γ, Σ . Il n'est pas utile que la coercion de la variable étende l'environnement, car le seul objet dans sa portée est la variable d'eta-expansion x . Pour plus de détails, voir la description de la règle `COERETAARR` de la Figure 4.11.

$$\begin{array}{c}
\text{SOUS-TYPAGE} \\
\frac{\Gamma \vdash \tau \text{ type} \quad \Gamma \vdash \sigma' \triangleright \sigma}{\Gamma \vdash \tau \triangleright \tau' \quad \Gamma \vdash \sigma' \triangleright \sigma} \\
\Gamma \vdash \tau' \rightarrow \sigma' \triangleright \tau \rightarrow \sigma
\end{array}
\qquad
\begin{array}{c}
\text{COERCION} \\
\frac{\Gamma \vdash \tau \text{ type} \quad \Gamma, \Sigma \vdash (\emptyset \vdash \tau) \triangleright \tau' \quad \Gamma \vdash (\Sigma \vdash \sigma') \triangleright \sigma}{\Gamma \vdash (\Sigma \vdash \tau' \rightarrow \sigma') \triangleright \tau \rightarrow \sigma}
\end{array}$$

Abstraction de coercions

L'abstraction de coercions est une abstraction effaçable, tout comme l'abstraction de types, autrement dit le polymorphisme. Cette thèse définit l'abstraction de coercions à l'aide de l'abstraction de types, en contraignant l'ensemble des types avec lesquels il sera possible d'instancier l'abstraction. L'exemple le plus simple est le cas de Système $F_{<}$: qui abstrait sur des types plus petits qu'une borne supérieur. Lorsque l'on veut instancier une telle abstraction, il faut montrer que le type argument est plus petit que la borne. Un exemple plus complet est le cas de ML avec contraintes qui abstrait sur des types vérifiant un ensemble de contraintes.

Le polymorphisme contraint n'est pas nécessairement cohérent : certaines abstractions de types ne sont pas instanciables. Ce cas de figure est discuté dans la Section Polymorphisme incohérent. Les abstractions de types cohérentes peuvent être classées dans trois catégories suivant la façon dont le type abstrait est contraint. Ces cas de figure sont présentés dans la Section Polymorphisme cohérent.

Polymorphisme cohérent

Le polymorphisme cohérent est la version la plus courante du polymorphisme contraint et la seule version effaçable. La plupart des systèmes de types avec polymorphisme, comme Système F, MLF ou Système $F_{<}$, ont un polymorphisme cohérent par construction : il n'est pas possible syntaxiquement d'écrire une abstraction de types incohérente. On peut différencier trois cas de polymorphisme cohérent : le polymorphisme non contraint, le polymorphisme borné, et le polymorphisme contraint. Chaque version contient les versions précédentes.

Polymorphisme non contraint

Le polymorphisme non contraint est celui que l'on trouve dans Système F. On peut instancier un type polymorphisme avec n'importe quel type du langage. Les abstractions de types sont alors cohérentes, car on peut les instancier avec top ou bottom par exemple. Le polymorphisme non contraint ne permet pas d'abstraire sur des coercions par définition.

Polymorphisme borné

La contrainte de polymorphisme la plus simple qui permette d'abstraire sur des coercions est la contrainte par une borne unique. Cette borne peut être supérieure ou inférieure. Par exemple, le polymorphisme de Système $F_{<}$ est borné par une borne supérieur, alors que celui

de MLF possède une borne inférieure. Dans notre système F_t^p (Chapitre 4), le polymorphisme peut être, soit non contraint comme dans Système F, soit borné par une borne supérieure comme dans Système $F_{<}$, soit borne par une borne inférieure comme dans MLF.

Tout comme pour le polymorphisme non contraint, le polymorphisme borné est toujours cohérent. Il est toujours possible d’instancier l’abstraction de types par le type extrême qui convient : top pour les bornes inférieure et bottom pour les bornes supérieure.

Il est possible qu’une borne fasse référence à la variable de type qu’elle contraint, on parle alors de borne récursive. Les bornes récursives n’ajoutent aucune difficulté et sont disponibles dans Système F_t^p .

Polymorphisme contraint

La contrainte maximale pour le polymorphisme cohérent est la cohérence de contrainte, à savoir qu’elle soit satisfiable au moins une fois. Ce type de polymorphisme est par exemple présent dans ML avec contraintes et dans notre système F_{cc} (Chapitre 5). Dans Système F_{cc} , on étend la syntaxe des kinds de Système F, qui ne contient alors que la kind étoile, avec une notion de kind contrainte qui rejette les types ne satisfiant pas la contrainte. Il est alors possible d’utiliser des contraintes de coercions pour abstraire sur des coercions.

Polymorphisme incohérent

Lorsqu’une abstraction de types est incohérente, la réduction doit être bloquée dans son corps, sans quoi le système de types accepterait des programmes incorrects. Un contre-exemple serait de prendre la négation booléenne d’un entier sous une abstraction qui nous garanti qu’un entier peut être vu comme un booléen. La réduction doit donc être faible sur les abstractions de types incohérentes, qui deviennent donc non-effaçables à l’inverse de leurs homologues cohérentes. Le langage noyau de GHC, à savoir Système FC, et notre langage Système F_{cc} disposent de polymorphisme incohérent. L’incohérence est une caractéristique nécessaire pour implémenter les GADTs (Section 5.5).

Types récursifs

Dans cette thèse, les types récursifs sont présentés comme des types iso-récursifs, mais disposent de l’expressivité des types equi-récursifs. Ils sont présentés comme des types iso-récursifs, car l’introduction et l’élimination des types récursifs sont des coercions : un type récursif est isomorphe (et non égal) à son dépliement. Cependant, ils ont l’expressivité des types equi-récursifs car ils sont effaçables, que il existe une coercion de congruence pour tous les types et donc tous les contextes de types, et enfin car ils disposent des règles de conversion des types equi-récursifs (Section 5.4.4). Le premier point vient du fait que les coercions sont effaçables. Le second point vient du fait que l’on fournit une règle d’eta-expansion pour tous les types (Section Eta-expansion) et que la congruence des types effaçables est dérivable. Enfin, le dernier point est possible car Système F_{cc} dispose d’un mécanisme de coinduction : il est possible de prouver une coercion avec elle-même pourvu qu’on n’y fasse référence que sous des types de calcul.

Sémantique

Une contribution annexe de cette thèse est l'adaptation à la réduction forte des méthodes de preuve par step-index en s'inspirant des preuves par candidats de réductibilité. Les preuves de correction peuvent être faites de différentes manières : syntaxiquement par préservation et progrès, sémantiquement, ou d'autres solutions. Nous avons tenté de faire une preuve syntaxique de correction pour Système F_{cc} (Section 4.6.2), mais cela complique trop le système. Nous avons donc opté pour une preuve sémantique. Cependant, la présence de types récursifs casse la preuve usuelle par candidats de réductibilité, qui repose sur la terminaison des termes. Comme la correction ne dépend que d'un nombre fini de pas de réduction, une solution est donc de considérer toutes les réductions finies des termes, qui terminent par définition. C'est le principe des step-index : les termes sont munis d'un index qui compte le nombre de pas de réductions restants. Cependant, cette technique ne permet pas d'étudier la réduction forte. Notre solution est de remplacer la notion d'index global de terminaison par une notion de terme indexé : chaque nœud du terme a son propre index de contrôle de la réduction (Section 5.2).

Contents

1	Introduction	17
2	The λ-calculus	21
2.1	Terminology	21
2.2	Syntax	22
2.3	Reduction rules	23
2.4	Encodings	26
2.4.1	Booleans	26
2.4.2	Pairs	27
2.4.3	Sums	27
2.5	Properties	28
2.5.1	Confluence	28
2.5.2	Curryfication	29
2.5.3	Soundness	29
2.5.4	Termination	29
I	Type Systems as Usual	31
3	Existing Type Systems	33
3.1	The STLC	34
3.1.1	Definition	34
3.1.2	Properties	37
3.2	System F	39
3.2.1	Definition	39
3.2.2	Properties	41
3.3	System F_{rec}	44
3.3.1	Definition	44
3.3.2	Properties	47
3.4	System F_{η}	48
3.4.1	Definition	48
3.4.2	Properties	52
3.5	MLF	53
3.5.1	Definition	53
3.5.2	Properties	58
3.6	System $F_{<}$	58

3.6.1	Definition	58
3.6.2	Properties	62
3.7	Constraint ML	63
3.8	Existing Coercions	64
II Type Systems as Coercions		67
4	An explicit calculus of coercions: System F_t^p	69
4.1	Base system	71
4.2	Features	74
4.2.1	Polymorphism	75
4.2.2	Eta-expansion	76
4.2.3	Bottom	77
4.2.4	Top	78
4.2.5	Lower Bounded polymorphism	78
4.2.6	Upper Bounded polymorphism	79
4.3	System F_t^p	80
4.4	Properties	88
4.4.1	Implicit vs. Explicit version	89
4.4.2	Termination	90
4.4.3	Confluence	92
4.4.4	Bisimulation	93
4.4.5	Soundness	95
4.5	Expressivity	99
4.5.1	System F	99
4.5.2	System F_η	100
4.5.3	MLF	102
4.5.4	System $F_{<}$	103
4.5.5	Summary	104
4.6	Beyond parametric coercion abstraction	104
4.6.1	Unrestricted coercion abstraction	105
4.6.2	Push	106
5	An implicit calculus of coercions: System F_{cc}	111
5.1	Definition	111
5.2	Semantics	120
5.2.1	The Indexed Calculus	121
5.2.2	Bisimulation	123
5.2.3	Semantic types	124
5.2.4	Simple types	125
5.2.5	Intersection types	127
5.2.6	Recursive types	127
5.2.7	Semantic judgment	128
5.3	Soundness	130
5.4	Expressivity	135
5.4.1	Surface notations	135

5.4.2	System F_t^p	137
5.4.3	Constraint ML	139
5.4.4	Recursive coercions	140
5.5	Incoherent Polymorphism	141
5.6	Coq formalization	146
5.7	Discussion about the explicit version	148
6	Discussions	151
6.1	Extensions	151
6.1.1	Data types	151
6.1.2	Existentials	151
6.1.3	Type-level functions	152
6.1.4	Recursive types at arbitrary kinds	153
6.1.5	Non-erasable coercions	154
6.1.6	First-class coercions	155
6.1.7	Dependent types	156
6.1.8	Kind coercions	156
6.1.9	Intersection types	157
6.1.10	Semantic consistency	157
6.1.11	Environment actions	158
6.1.12	Coercion reduction	158
6.1.13	Side effects	159
6.1.14	Dead code	159
6.2	Related work	160
6.2.1	System $F_{<}$	160
6.2.2	System FC	161
6.2.3	Implicit Coercions	161
6.2.4	Step-indices	162
6.3	Applications	162
7	Conclusion	165

Chapter 1

Introduction

A *type system* has exactly one underlying programming language. It classifies its programs in order to reject those that may go wrong. Sometimes, it also require programs to terminate. Type systems do so by defining abstractions of program behaviors, which are verified by the type checker at compile-time. Thus, type systems also play a role for code documentation. Here, we restrict our attention to type systems for functional programming languages. A programming language is functional if its programs can take a function as argument or return a function as a result.

Type systems usually come with some *inference* mechanism. Inference permits the programmer to write fewer typing annotations. OCaml and Haskell have a powerful inference mechanism, while Coq requires more guidance from the user. This difference comes from the fact that Coq has a more expressive type system than these two languages. We call a *surface type system* a type system defined for the programmer. Surface type systems usually have a powerful inference mechanism, which makes them easier to program with. We call a *kernel type system* a type system defined for the language designer. They are usually more concise in definition and built to have fewer, more fundamental features. Sometimes the kernel type system of a language is simply the fully-explicit version of the surface type system, in which case the differences are mainly syntactical. But sometimes there is some desugaring and term must be elaborated from the surface type system to the kernel type system. In this case, the kernel type system usually has fewer constructs, which makes it shorter to formalize. However, the kernel type system has to be more expressive than the surface one, which on the opposite usually makes it harder to formalize than the surface language. For instance, the GHC implementation of Haskell has a surface type system, Haskell with extensions (GADTs, type families, type classes, etc.), and a kernel type system, called FC (System F with equality coercions and coercion abstraction). Type classes, GADTs, and type families are surface type system features, and do not need to be proven sound in FC. However, FC has to deal with equality coercions and coercion abstraction. It is less fastidious to prove soundness in FC, although it may be more involved conceptually.

The untyped λ -calculus with constants (pairs, integers, etc.) is the simplest functional programming language. It has by definition no type system, and is thus a programming language in this sense. The λ -calculus is the underlying language of type systems such as Coq, Haskell, and OCaml for example. Each of them has its own practical implementation and its own type system. Although their underlying programming language share the λ -calculus, none of them is exactly the λ -calculus. For example, Haskell and OCaml provide

side effects and use a particular reduction strategy. Both use weak reduction, but Haskell uses a call-by-need strategy while OCaml uses a call-by-value strategy. Coq is *pure* (without side effects), uses strong reduction, and its type system also ensures strong normalization.

Although these type systems have their own particularities, they still share the λ -calculus for their core language. As such, it is interesting to study the possibility for the λ -calculus to have a type system that would ensure the soundness and strong normalization of other more exciting type systems. The soundness and normalization for strong reduction imply the soundness and normalization for all usual strategies. As a consequence, studying soundness and normalization in a strong reduction setting is enough. Studying only the pure λ -calculus (without side effects) still gives a result for languages with side effects when the type system keeps track of side effects with a monadic encoding as it is the case in Haskell. Although dependent type systems are a desire, their design is still a research topic in itself, so we do not include them here.

To define such general type system for the λ -calculus with strong reduction, we generalize existing ideas: several type systems use subtyping, containments, or coercions to express some of their features. For instance, System $F_{<}$ uses subtyping to retype a term or to constrain type abstraction. GHC uses equality coercions to implement GADTs and type families from its surface type system to its kernel type system.

In this document we introduce a framework for defining and studying type systems based on *typing coercions*. We define typing coercions (which we simply call coercions) as composable and erasable typing transformations. We instantiate the framework on two type systems where all features are expressed as coercions on top of the STLC. In particular, we make a distinction between *computational types* (usually called simple types) and *erasable types*. Computational types are related to the reduction and terms, while erasable types are related to typing. Coercions introduce and eliminate erasable types, while they can only provide congruence rules for computational types. Congruence rules for erasable types are derivable from their introduction and elimination rules. Moreover, we exhibit and explain one fundamental feature behind all of these type systems: coercion abstraction.

The first type system we define, called System F_t^P , subsumes a few existing functional type systems: System $F_{<}$ [5], MLF [19], and System F_η [25]. System $F_{<}$ features upper bounded polymorphism and arrow congruence. MLF extends System F with lower bounded polymorphism to feature a complete inference mechanism as long as term variables used polymorphically are annotated. System F_η extends System F with η -expansion: a term is well-typed in System F_η if one of its η -expansion is well-typed in System F. The main particularity of System F_η , compared to usual type systems with subtyping, is the distributivity rule $\forall\alpha(\tau \rightarrow \sigma) \leq (\forall\alpha\tau) \rightarrow \forall\alpha\sigma$. Our type system F_t^P combines the following coercion features: polymorphism, lower bounded polymorphism, upper bounded polymorphism, and η -expansion. These features seen as coercions permit to derive the distributivity rule and the polymorphism congruence rule of System F_η , the upper bounded polymorphism congruence rule of the most expressive version of System $F_{<}$, and the lower bounded polymorphism congruence rules of MLF. Lower and upper bounded polymorphism can be seen as restricted forms of coercion abstraction. We call this restriction parametric because the abstract coercion has to be parametric on either its argument type (upper bounded polymorphism) or its result type (lower bounded polymorphism). This restriction ensures that coercion abstractions are coherent, *i.e.* that they do not introduce inconsistencies to the current environment.

The second type system we define, called System F_{cc} , subsumes additional type systems: System F_{rec} and Constraint ML [27]. System F_{rec} extends System F with general equi-recursive

types. **Constraint ML** extends **ML** (prenex polymorphism with complete inference) with sub-typing and still provides type inference. As a kernel type system the main feature of **Constraint ML** is constraint abstraction, which is similar to coercion abstraction. System F_{cc} combines coherent polymorphism, η -expansion, and recursive types as coercions. It can also constrain a set of types to the types satisfying a proposition. Since coercions are one sort of propositions, this constrained kind feature and polymorphism permit unrestricted coercion abstraction.

Although System F_{cc} subsumes System F_t^p , it is not better on every point. System F_t^p has an explicit version where explicit terms and explicit coercions capture the essence of typing derivations. The proof of subject reduction is thus done by reduction of the explicit term witnessing the derivation. System F_{cc} has only an implicit version without a subject reduction property, even though the type system is shown to be sound. This difference comes from the fact that an explicit version of System F_{cc} with subject reduction would need coercion decomposition and a push reduction rule. When we have a coercion from $\tau' \rightarrow \sigma'$ to $\tau \rightarrow \sigma$ in-between a redex, we have to decompose it into two coercions in order to progress: one coercion from σ' to σ and one from τ to τ' . The main difficulty with this extension is the proof of consistency. We managed to show one side of the decomposition, however we do not know whether the other side holds in our semantics. Moving polymorphism from the coercion language to the term language may be a solution to restore subject reduction at the cost of losing deep type and coercion abstraction and instantiation.

Both System F_t^p and System F_{cc} can be seen as kernel type systems. The type systems they subsume can then be seen as associated surface type systems. For instance, type systems with inference such as **MLF** or **Constraint ML** are surface type systems for System F_{cc} .

In addition to the definition of the coercion framework, the main contributions of this thesis are the study of coercion abstraction, the definition of a general η -expansion rule, and a step-indexed semantics for strong reduction. All this work is done in a strong reduction setting in order for the soundness result to be applicable to all usual strategies (those included in strong reduction, like weak reduction and its call-by-value and call-by-need restrictions). Stating the soundness property in a strong reduction setting also permits to capture more of the information contained in the type system. For instance, when a term admits a typing derivation, it is usually sound under term abstractions while the soundness property in weak reduction does not imply it, since reduction may not look under lambdas. This thesis, by defining a common framework describing several type systems, permits to easily compare the features of these systems. A common underlying programming language is highly valuable when comparing two type systems. It is more interesting to say that a type system subsumes another one, than to say that it encodes the other one. In the first case, the same term can be used from one type system to the other. Finally, a first contribution of System F_t^p was to prove the strong normalization of **MLF**.

My thesis is that defining erasable type system features as coercions, *i.e.* composable inclusions between invariants, naturally gives the most out of a type system feature and its combination with other coercion features. For instance, combining η -expansion and upper bounded polymorphism naturally gives a type system more expressive than the most expressive version of System $F_{<}$. Once the type system is defined as coercions, adding coercion abstraction as a coercion extends the expressivity even further. Finally, the study of composable type systems in a strong reduction setting gives more information than with a particular strategy. A composable type system does not distinguish between top-level expressions and expressions under arbitrary environments. In particular, it does not have specialized rules for typing top-level closed terms. As a consequence, results for well-typed closed terms are

likely to also hold for well-typed open terms. This is, for example, the case for soundness and normalization.

Chapter 2

The λ -calculus

On the one hand, programming languages such as C, C++, or Java, have functions defined at top-level. In other words, functions are statements and not expressions. For instance, it is not possible to pass a function as argument to another function without using pointers or wrapper classes. A typical need for this kind of feature is the *map function*. The map function takes a transformation function and a collection as arguments. It returns a collection similar to the original collection where all elements have been transformed using the transformation function.

On the other hand, *functional programming languages* use the paradigm of first-class functions. In other words, functions are expressions and not statements. It is possible to pass a function as argument to another function and thus it is possible to define a map function for all kinds of collections. It is also possible to define an anonymous function for local use, for example, inside an arithmetical expression.

The λ -calculus is a functional programming language invented by Alonzo Church. It is at the same time, simple to explain and reason about, and able to describe high-level programming. It is actually used as the foundation of several general-purpose programming languages, such as OCaml [26] or Haskell [17]. These languages are also functional programming languages.

All type systems, existing or new, exposed in this document are based on the λ -calculus. This chapter describes the syntax, the reduction rules, and some properties of the λ -calculus. We only present the properties we refer to in subsequent chapters, since this calculus satisfies numerous properties. Moreover, relying on fewer properties shows that our framework for type systems applies to more calculi. For instance, our results do not rely on the determinism, confluence, or standardization of the calculus (see 6.1.2 where standardization may become necessary).

2.1 Terminology

In the λ -calculus, programs are called *terms*. So I use programs to refer to programs of arbitrary languages, and terms to refer to programs of the λ -calculus. As in mathematics we use the letter x for variables. We use the letters a and b for terms. The notion of execution is generalized to open terms and called *reduction*. The result of a good execution is called a *value*. The definition of a function is called the *body* of the function.

2.2 Syntax

The λ -calculus main characteristic is to place functions at the center of the language by making them usual programs: functions are terms. In particular in the pure λ -calculus all terms are functions. However the λ -calculus may be easily extended with additional terms. For instance we can add pairs, tags, or user-defined data-types.

What is a function and how do we use it? A function in programming languages can be seen, as in mathematics, as a relation between an input and an output, with the property that each input is in relation with exactly one output. To use a function we need to provide an input. This input will be substituted in the function body for the input variable.

For instance, f defined by $f(x) = x + 1$ is a mathematical function between natural numbers. The input natural number x is in relation with a unique output which in this case is its successor. To use it we need to say with which natural number we want to evaluate it. To do so, we replace x by the chosen natural number. For instance, $f(1) = 2$ and $f(2) = 3$.

But we can also see functions, as in programming languages, as subroutines: a function is a program where some parts are not yet defined. We say that the function abstracts over these parts. To use such a function we supply it with concrete parts for the abstract parts. It is actually enough to study functions that abstract over exactly one term using currying, see 2.5.2.

For instance, f defined by $f(x) = x + 1$ is a function between expressions. The input expression x is in relation with a unique output which in this case is the expression $x + 1$. To use it we need to say with which expression we want to substitute. To do so, we replace x by the chosen expression. For instance, $f(3 \times 4) = 3 \times 4 + 1$ and $f(f(1)) = f(1) + 1$.

According to these definitions, the three notions we need to create and use functions are:

- The notion of term abstraction which we write $\lambda x a$ where a is the definition, called body, of the function and x is its abstract part. For instance, $\lambda x x + 1$ would be the definition of the function returning $x + 1$ for x .
- The notion of term application which we write $a b$ where a is the function and b is the argument. See below for more details.
- The notion of term variable which we write x . It designates the places where the concrete argument will be substituted.

It may at first seem strange why application is written $a b$ instead of $x(a)$ as it would be the case in mathematics or other programming languages, where x is the name of the function and a its argument. In non-functional programming languages and in mathematics, functions are first named, and then used. It is not possible to use a function anonymously, without giving it a name. In the λ -calculus, since functions are terms and not only definitions, they don't have names by default. We will see in the next paragraph how we can give names to terms. Since names (variables) are also terms, when we say that applications take a term to be used as a function we allow both names and functions to be used. See Section 2.3 to see that we can use even more terms as functions. Finally, we do not need parentheses by default around the argument. It is only required to add parentheses when something is ambiguous. Notice that we define a notion of construct precedence in order to have fewer ambiguous terms (see the paragraph about construct precedence).

x, y	Term variables
$a, b ::= x \mid \lambda x a \mid a a \mid \langle a, a \rangle \mid \text{fst } a \mid \text{snd } a$	Terms

Figure 2.1: Syntax of the λ -calculus

Giving names In mathematics and programming languages we can name objects and functions. The same effect is already possible in the pure λ -calculus using only the three constructions (abstraction, application, and variable) we defined. The construction **let** $x = b$ **in** a can be defined as syntactic sugar for $(\lambda x a) b$, the application of the argument term b to the function term $\lambda x a$. To understand why it gives b the name x in a , let's look at its definition. We apply the argument b to the function abstracting over x in a , so b is the concrete term for the abstract term x in a , which is exactly what we wanted.

Construct precedence Terms can be given either by their abstract syntax tree or as text. In this latter case some ambiguity might appear. Take the following term $a_1 a_2 a_3$ for instance. It is not clear whether it is a_1 applied to a_2 and a_3 , namely $(a_1 a_2) a_3$, or a_1 applied to a_2 applied to a_3 , namely $a_1 (a_2 a_3)$. To resolve this ambiguity we let the application constructor be left associative: the first form is the right one in the absence of parentheses. Similarly the term $\lambda x a b$ can either be $\lambda x (a b)$ or $(\lambda x a) b$. To resolve this last ambiguity we give precedence to the application: the first form is the right one in the absence of parentheses.

Pairs The soundness property defined in Section 2.5.3 is only meaningful when we extend the pure λ -calculus with other constructors, since otherwise it trivially holds. To keep things simple, we add pairs to the λ -calculus, but other constructors could be added as well: records, variants, etc. To construct the pair of a and b we write $\langle a, b \rangle$. To extract the first component from a pair a we write $\text{fst } a$ and to extract its second component we write $\text{snd } a$.

Summary All these definitions are summarized in Figure 2.1. We write x and y for variables and a and b for terms. The set of terms is inductively defined. Terms can be either variables x , abstractions $\lambda x a$, applications $a a$, pairs $\langle a, a \rangle$, or projections $\text{fst } a$ or $\text{snd } a$.

2.3 Reduction rules

Programs can be executed. This either returns a result or loops indefinitely. The same mechanism applies for the λ -calculus, even if we do not study side-effects. The notion of execution is called *reduction*. Reduction is a relation between terms, written $a \rightsquigarrow b$, which means that the term a reduces on the term b in one step. This relation is not a function, which means reduction is a priori non-deterministic.

In mathematics expressions are evaluated. For instance, to evaluate $f(2)$, we look for the definition of f , say $f(x) = x + 1$, and we substitute x in the definition by its value. We get $2 + 1$ which in turn evaluates to 3 . In programming languages programs are executed. Most of the time it is a list of instructions to be executed sequentially. When an instruction is a call to a subroutine, the list of instructions in the subroutines are prepended to the remaining instructions. The same mechanism applies to the λ -calculus. When a function is applied to an argument as in $(\lambda x a) b$, it can reduce the body of the function, namely a , after the variable x has been substituted by the argument b . We use the notation $a[x/b]$ to denote the substitution

$$\begin{array}{c}
\text{REDCTX} \\
\frac{a \rightsquigarrow b}{E[a] \rightsquigarrow E[b]}
\end{array}
\qquad
\begin{array}{c}
\text{REDAPP} \\
(\lambda x a) b \rightsquigarrow a[x/b]
\end{array}
\qquad
\begin{array}{c}
\text{REDFST} \\
\text{fst } \langle a, b \rangle \rightsquigarrow a
\end{array}
\qquad
\begin{array}{c}
\text{REDSND} \\
\text{snd } \langle a, b \rangle \rightsquigarrow b
\end{array}$$

Figure 2.2: Reduction relation

$$\begin{array}{ll}
E ::= \lambda x \square \mid \square a \mid a \square \mid \langle \square, a \rangle \mid \langle a, \square \rangle \mid \text{fst } \square \mid \text{snd } \square & \text{Evaluation contexts} \\
p ::= x \mid p v \mid \text{fst } p \mid \text{snd } p & \text{Prevalues} \\
v ::= p \mid \lambda x v \mid \langle v, v \rangle & \text{Values} \\
r ::= E[r] \mid \langle a, a \rangle a \mid \text{fst } (\lambda x a) \mid \text{snd } (\lambda x a) & \text{Errors}
\end{array}$$

Figure 2.3: Notations

of x by b in a . This reduction rule can then be written $(\lambda x a) b \rightsquigarrow a[x/b]$. We call this rule **REDAPP** and give it in Figure 2.2.

Notice that the substitution to make sense, has to be capture avoiding. This means that if the argument b contains free variables (variables not bounded by an abstraction), they should remain free after the substitution $a[x/b]$. Concretely, if a is $\lambda y y x$ and b is y , the substitution $(\lambda y y x)[x/y]$ should return $\lambda y_1 y_1 y$ and not $\lambda y y y$. We renamed the variable y with the fresh variable y_1 in a . This operation is called α -conversion. It does not modify the meaning of the term and is always possible. We say that $\lambda y y x$ is α -equivalent to $\lambda y_1 y_1 x$. In the following we treat terms up to α -equivalence.

We also have two rules for pairs: one for the first projection of a pair and one for the second one. The first projection of a pair, namely **fst** $\langle a, b \rangle$, reduces on the first component of the pair, namely a . We write this reduction rule **fst** $\langle a, b \rangle \rightsquigarrow a$ and name it **REDFST**. We have a similar rule for the second projection.

Finally, we want reduction to occur anywhere in the term. This is called strong reduction. Properties that hold for strong reduction also holds for reduction considering less evaluation contexts. In particular the soundness property for any smaller reduction relation is a consequence of the soundness property for strong reduction. Evaluation contexts are defined in Figure 2.3. They are terms where exactly one subterm was replaced by a hole, written \square .

We can now give all reduction rules in Figure 2.2. Rule **REDCTX** is the context reduction rule. If a subterm a reduces to b under the evaluation context E , then the term $E[a]$, where the hole of the evaluation context E was replaced by the subterm a , reduces to $E[b]$. We write $a \rightsquigarrow^+ a$ the transitive closure of the reduction relation $a \rightsquigarrow a$. And we write $a \rightsquigarrow^* a$ its reflexive and transitive closure. We call a redex the left-hand side of rules **REDAPP**, **REDFST**, and **REDSND**.

In the absence of side effects, the goal of reduction is to reach a value. For instance, in mathematics, when we evaluate an expression, we want the evaluation to terminate and return a result. We will talk about termination in Section 2.5.4 but we can already talk about results in the λ -calculus. When the reduction of a term terminates, it reaches an irreducible term. Irreducible terms are terms that cannot reduce. Irreducible terms are results of reduction, but they are not necessarily values. Some irreducible terms are values, some are errors.

Errors can either be immediate errors or errors under an evaluation context. An immediate error looks like a redex but it is not a redex. To define what looks like a redex we need to define

the notions of constructors and destructors. A constructor is either an abstraction or a pair, while a destructor is either an application context or a projection context. Redexes are usually a destructor applied to an associated constructor. For instance, the redex $(\lambda x a) b$ applies the constructor $\lambda x a$ to the destructor $\square b$. The term $\langle a_1, a_2 \rangle b$ looks like a redex because it applies the constructor $\langle a_1, a_2 \rangle$ to the destructor $\square b$. However pairs are not associated to applications and the given term cannot be reduced: it is an immediate error. Similarly, abstractions are not associated to projections and $\mathbf{fst} (\lambda x a)$ and $\mathbf{snd} (\lambda x a)$ are immediate errors too. We write r for errors and Ω for the set of errors. Errors are given in Figure 2.3. We call valid, a term which is not an error. We write \mathcal{V} for the set of valid terms. It is by definition the complement of Ω .

Values are valid irreducible terms. Because they are valid they should not contain errors. And because they are irreducible, they should not contain redexes or they would not be irreducible. As a consequence, values are terms that do not contain errors or redexes. Because errors and redexes are constructors applied to destructors, we isolate the values that do not start with a constructor and call them prevalues. We define the set of values and prevalues inductively in Figure 2.3. We write p for prevalues and v for values. We call neutrals terms, the terms that do not start with a constructor and head normal forms those that do. These two sets of terms are complement of each other. So prevalues are neutral values.

Prevalues contain variables, because they are valid irreducible neutral terms. Prevalues also contain applications of prevalues to values $p v$, because both p and v do not contain errors or redexes by induction, and the application itself is not an error or a redex as p is neutral. Finally, prevalues contain projections of prevalues $\mathbf{fst} p$ and $\mathbf{snd} p$, because p does not contain redexes of any sort and the projection itself is not an error or a redex as p is neutral. Values contain prevalues by definition and constructors applied to values: $\lambda x v$ and $\langle v, v \rangle$. We can unfold these definitions and say that values are a series of constructors followed by a series of destructors applied to variables.

Examples Let's consider the following example: $\mathbf{let} y = \lambda x x \mathbf{in} y y$. We give the name y to the term $\lambda x x$, and then we apply it to itself. In order to reduce this term we first need to unfold the \mathbf{let} sugar. We look at its definition. The expression $\mathbf{let} x = b \mathbf{in} a$ stands for $(\lambda x a) b$, which reduces to $a[x/b]$. So we have that $\mathbf{let} x = b \mathbf{in} a$ reduces to $a[x/b]$. So our example reduces to $(\lambda x x) (\lambda x x)$ which itself reduces to $\lambda x x$. We observe that $\lambda x x$ always returns its argument, we call this term the identity function and use \mathbf{id} as sugar for it:

$$\mathbf{id} \triangleq \lambda x x$$

Another interesting term is the looping combinator \mathbf{omega} defined as sugar:

$$\mathbf{omega} \triangleq \mathbf{let} y = \lambda x x x \mathbf{in} y y$$

Let's see how it reduces. We first unfold its definition, then unfold the \mathbf{let} -definition. We can now reduce it by rule REDAPP. We α -rename the first abstraction from x to y . We can now fold the \mathbf{let} -definition and \mathbf{omega} -definition:

$$\begin{aligned} \mathbf{omega} &\triangleq \mathbf{let} y = \lambda x x x \mathbf{in} y y \triangleq (\lambda y y y) (\lambda x x x) \\ &\rightsquigarrow (\lambda x x x) (\lambda x x x) \\ &\stackrel{\alpha}{=} (\lambda y y y) (\lambda x x x) \triangleq \mathbf{let} y = \lambda x x x \mathbf{in} y y \triangleq \mathbf{omega} \end{aligned}$$

As a conclusion we have $\mathbf{omega} \rightsquigarrow \mathbf{omega}$, so the looping combinator reduces on itself in one step and this is the only reduction it can do. This is an example of a non-terminating term: it has no result (good or bad). But it never reaches an error either.

Strategies It is possible to choose a subrelation of the reduction relation that would be a partial function. This is the case for most real-life programming languages. Almost all programming languages use weak reduction, which consists in removing the abstraction context $\lambda x \square$. They also reduce pairs from left to right by modifying the evaluation context $\langle a, \square \rangle$ to $\langle v, \square \rangle$. Then most programming languages use what is called a call-by-value strategy, while some languages use a call-by-name (or its call-by-need optimization) strategy. Each time evaluation contexts or the reduction relation are modified, the set of values has to be changed too.

In call-by-value we modify the evaluation context $a \square$ to $(\lambda x a) \square$. In other words we only reduce the argument of an application if its function is an abstraction. We also restrain rule REDAPP to values, which is why we call this strategy call-by-value. The rule becomes $(\lambda x a) v \rightsquigarrow a[x/v]$. Notice that the set of values has changed. It contains abstractions $\lambda x a$, and pairs of values $\langle v, v \rangle$.

In call-by-name we remove the evaluation context $a \square$. The set of values is the same as those of call-by-value.

2.4 Encodings

Although we can extend the pure λ -calculus with pairs, naturals, booleans, etc., we can already reason with them with encodings. These encodings are not as fast as their analog extensions (the typical example is for natural subtraction), but they can simulate the computations. Another difference with the pure λ -calculus is that extensions of it may contain errors and terms that reduce on errors, while the pure λ -calculus contain no errors. This is why we study the λ -calculus with pairs instead of the pure λ -calculus, when we are interested in the soundness property. The pure λ -calculus contain no errors because there is only one sort of constructor and one sort of destructor, and they are associated.

2.4.1 Booleans

Booleans come with two constructors, namely **true** and **false**, and one destructor, namely the if-statement, written **if a then a else a** . In the expression **if b then a_1 else a_2** , we say that b is the conditional and it should evaluate to a boolean. The term a_1 is the term to evaluate if the conditional is true, while a_2 is evaluated when the conditional is false. The two reduction rules are thus:

$$\begin{array}{ll} \text{REDTRUE} & \text{REDFALSE} \\ \text{if true then } a \text{ else } b \rightsquigarrow a & \text{if false then } a \text{ else } b \rightsquigarrow b \end{array}$$

We can encode these three constructs in the pure λ -calculus as follows:

- $\mathbf{true} \triangleq \lambda x \lambda y x$
- $\mathbf{false} \triangleq \lambda x \lambda y y$
- $\mathbf{if } b \text{ then } a_1 \text{ else } a_2 \triangleq b a_1 a_2$

Let's verify that the two expected reduction rules work:

- if true then a else $b \triangleq (\lambda x \lambda y x) a b \rightsquigarrow (\lambda y a) b \rightsquigarrow a$
- if false then a else $b \triangleq (\lambda x \lambda y y) a b \rightsquigarrow (\lambda y y) b \rightsquigarrow b$

We can now define additional functions:

- not $\triangleq \lambda x$ if x then false else true
- and $\triangleq \lambda x \lambda y$ if x then y else false
- or $\triangleq \lambda x \lambda y$ if x then true else y

2.4.2 Pairs

Pairs can also be encoded in the pure λ -calculus as follows:

- $\langle a, b \rangle \triangleq \lambda y$ if y then a else b
- $\text{fst } a \triangleq a \text{ true}$
- $\text{snd } a \triangleq a \text{ false}$

Let's verify that the two expected reduction rules work:

$$\begin{aligned}
 \text{fst } \langle a, b \rangle &\triangleq (\lambda x x \text{ true}) ((\lambda x_1 \lambda x_2 \lambda y \text{ if } y \text{ then } x_1 \text{ else } x_2) a b) \\
 &\rightsquigarrow (\lambda x_1 \lambda x_2 \lambda y \text{ if } y \text{ then } x_1 \text{ else } x_2) a b \text{ true} \\
 &\rightsquigarrow (\lambda x_2 \lambda y \text{ if } y \text{ then } a \text{ else } x_2) b \text{ true} \\
 &\rightsquigarrow (\lambda y \text{ if } y \text{ then } a \text{ else } b) \text{ true} \\
 &\rightsquigarrow \text{if true then } a \text{ else } b \\
 &\rightsquigarrow a
 \end{aligned}$$

$$\begin{aligned}
 \text{snd } \langle a, b \rangle &\triangleq (\lambda x x \text{ false}) ((\lambda x_1 \lambda x_2 \lambda y \text{ if } y \text{ then } x_1 \text{ else } x_2) a b) \\
 &\rightsquigarrow (\lambda x_1 \lambda x_2 \lambda y \text{ if } y \text{ then } x_1 \text{ else } x_2) a b \text{ false} \\
 &\rightsquigarrow (\lambda x_2 \lambda y \text{ if } y \text{ then } a \text{ else } x_2) b \text{ false} \\
 &\rightsquigarrow (\lambda y \text{ if } y \text{ then } a \text{ else } b) \text{ false} \\
 &\rightsquigarrow \text{if false then } a \text{ else } b \\
 &\rightsquigarrow b
 \end{aligned}$$

2.4.3 Sums

Sums come with two constructors, namely the injections, and one destructor, namely the `case` operator. The left injection is written `inl a` and the right injection is written `inr a`. The case operator is written `case a of {inl $x_1 \mapsto b_1$ | inr $x_2 \mapsto b_2$ }`. If the argument a is of the form `inl a_1` ,

then the first branch b_1 is evaluated with x_1 substituted with a_1 . But if a is of the form $\text{inr } a_2$, a similar evaluation happens with the second branch b_2 . There two reduction rules are thus:

$$\begin{array}{l} \text{REDINL} \\ \text{case inl } a \text{ of } \{\text{inl } x_1 \mapsto b_1 \mid \text{inr } x_2 \mapsto b_2\} \rightsquigarrow b_1[x_1/a] \end{array}$$

$$\begin{array}{l} \text{REDINR} \\ \text{case inr } a \text{ of } \{\text{inl } x_1 \mapsto b_1 \mid \text{inr } x_2 \mapsto b_2\} \rightsquigarrow b_2[x_2/a] \end{array}$$

We can encode these constructs in the pure λ -calculus as follows:

- $\text{inl } a \triangleq \lambda y_1 \lambda y_2 y_1 a$
- $\text{inr } a \triangleq \lambda y_1 \lambda y_2 y_2 a$
- $\text{case } a \text{ of } \{\text{inl } x_1 \mapsto b_1 \mid \text{inr } x_2 \mapsto b_2\} \triangleq a (\lambda x_1 b_1) (\lambda x_2 b_2)$

Let's verify that the two expected reduction rules work:

$$\begin{aligned} \text{case inl } a \text{ of } \{\text{inl } x_1 \mapsto b_1 \mid \text{inr } x_2 \mapsto b_2\} &\triangleq (\lambda y_1 \lambda y_2 y_1 a) (\lambda x_1 b_1) (\lambda x_2 b_2) \\ &\rightsquigarrow (\lambda y_2 (\lambda x_1 b_1) a) (\lambda x_2 b_2) \\ &\rightsquigarrow (\lambda x_1 b_1) a \\ &\rightsquigarrow b_1[x_1/a] \end{aligned}$$

$$\begin{aligned} \text{case inr } a \text{ of } \{\text{inl } x_1 \mapsto b_1 \mid \text{inr } x_2 \mapsto b_2\} &\triangleq (\lambda y_1 \lambda y_2 y_2 a) (\lambda x_1 b_1) (\lambda x_2 b_2) \\ &\rightsquigarrow (\lambda y_2 y_2 a) (\lambda x_2 b_2) \\ &\rightsquigarrow (\lambda x_2 b_2) a \\ &\rightsquigarrow b_2[x_2/a] \end{aligned}$$

2.5 Properties

We present some properties of the λ -calculus, like confluence and curryfication. These properties are for the whole programming language. On another side, some properties are about terms, and all terms of the λ -calculus do not satisfy these properties. For example, soundness and termination are not true for all terms. Notice however that the pure λ -calculus is sound: all its term are sound. This comes from the fact that there is only one sort of constructor and one sort of destructor, and they are associated. There is syntactically no errors in the pure λ -calculus.

2.5.1 Confluence

A reduction relation is confluent if any two reduction paths starting from the same term can be joined. More precisely if a reduces in zero or more steps to a_1 and if it also reduces in zero or more steps to a_2 , then there is a term b such that a_1 and a_2 reduce in zero or more steps to b .

Definition 1 (Confluence). *A relation $a \rightsquigarrow a$ is confluent if it satisfies the following property. If $a \rightsquigarrow^* a_1$ and $a \rightsquigarrow^* a_2$ hold, then there is b such that $a_1 \rightsquigarrow^* b$ and $a_2 \rightsquigarrow^* b$ hold.*

The λ -calculus with pairs is confluent. Other extensions of the λ -calculus are also confluent, but we focus on the λ -calculus extended with pairs in this document. The λ -calculus is also locally confluent. Local confluence is similar to confluence but its hypotheses do exactly one reduction step.

Definition 2 (Local confluence). *A relation $a \rightsquigarrow a$ is locally confluent if it satisfies the following property. If $a \rightsquigarrow a_1$ and $a \rightsquigarrow a_2$ hold, then there is b such that $a_1 \rightsquigarrow^* b$ and $a_2 \rightsquigarrow^* b$ hold.*

2.5.2 Curryfication

In the λ -calculus, functions take exactly one argument and return exactly one result. However using pairs, it is possible to take several arguments and return several results. For instance, a swap function that takes two arguments at the same time and returns them swapped, could be written:

$$\text{swap} \triangleq \lambda x \langle \text{snd } x, \text{fst } x \rangle$$

It is however more convenient to use functions over several arguments like $\lambda(x, y) \langle y, x \rangle$ although we still have to return a pair. Functions taking several arguments can also be curryfied. This means that instead of abstracting over a series of arguments, the term abstracts the first component of the series and returns an abstraction over the second component of the series, and so on. More precisely $\lambda(x_1, \dots, x_n) a$ can be curryfied as $\lambda x_1 \dots \lambda x_n a$ and reciprocally for uncurryfication. A similar mechanism applies for multi-application. The application of a series of argument $a(b_1, \dots, b_n)$ becomes $(a b_1) \dots b_n$. The parentheses are not necessary.

2.5.3 Soundness

Any powerful enough programming language allows somehow to write programs that go wrong. For instance, null pointer exceptions or segmentation faults are consequences of programs that went wrong. In the pure λ -calculus, programs cannot go wrong since no errors can appear syntactically. However when we extend it, errors may appear, which is why we added pairs.

A term a is sound if and only if none of its reduction paths yield to an error (or all its reduction paths yield to valid terms).

Definition 3 (Sound terms). *A term a is sound if $\forall b, a \rightsquigarrow^* b \Rightarrow b \in \mathcal{U}$.*

2.5.4 Termination

Any powerful enough programming language allows some kind of looping mechanism (while-loops, for-loops, recursive functions, etc.) hence non-termination. The λ -calculus does not differ on this point, for instance the **omega** term loops indefinitely on itself. When we use a program to implement an algorithm, we are usually interested in its termination to obtain the resulting value. So we are interested into the termination of its reduction paths.

A term a is strongly normalizing if all of its reduction paths are finite.

Part I

Type Systems as Usual

Chapter 3

Existing Type Systems

In this chapter, we describe some existing type systems: one type system per section. All these type systems will be unified and subsumed by our final type system in Chapter 5, namely System F_{cc} . Understanding all these type systems in detail is not necessary, however it is important to notice the differences in their definition, presentation, and set of features. At a first glance, it is not obvious that they may be unified and that the features of one type system are compatible with the features of another type system.

Since this chapter presents a series of existing type systems, it may be boring to read for readers familiar with these notions. Each type system can be skipped independently, whether the reader know them or not. The only new content, which may thus be of interested to experienced readers, is the explicit version of System F_η in Section 3.4. However this extension comes without surprises.

In the previous chapter we saw that for powerful enough programming languages, some interesting properties about terms, like soundness and termination, do not hold for every term. We want a way to know when they hold to avoid buggy programs or looping algorithms. The first approach is to mathematically prove the properties we are interested in each time we consider a term. If we write a lot of programs this can quickly become not tractable. Moreover, this is neither modular, nor resilient to code changes. The second approach would be to use more automation. One way would be to use an automatic prover and help it when needed. But another successful approach is to use type systems. Besides, as said in the introduction, type systems are also a good starting point for code documentation.

A type system is a small syntactic language which objects can be interpreted as proofs. The main object is the term judgment, which is usually of the form $\Gamma \vdash a : \tau$ and tells that a is sound and in some type systems that it also terminates. These objects can either be completely written by the programmer, partially inferred or totally inferred. In the present document we do not discuss the inference issue and only look at the type system once all inference is done. Said otherwise we only study kernel type systems.

Type systems can either be implicit or explicit. Explicit type systems have an additional syntactical class called explicit terms. They are partial typing derivations and contain all necessary annotations for the typing derivations to be unique under some conditions. The types are then unique for each initial environment. The explicit version simply gives a name to each rule and is thus uniquely defined by the implicit version. Moreover, explicit type systems define a notion of reduction over these explicit terms that bisimulates the term reduction, up to erasable steps. This reduction is directly used to show the subject reduction property by

$M, N ::= x \mid \lambda(x : \tau) M \mid M M \mid \langle M, M \rangle \mid \text{fst } M \mid \text{snd } M$	Explicit terms
$\tau, \sigma, \rho ::= \tau \rightarrow \tau \mid \tau \times \tau$	Types
$\Gamma ::= \emptyset \mid \Gamma, (x : \tau)$	Environments

Figure 3.1: STLC syntax

reduction of the typing derivation. This is a very strong property that all type systems do not necessarily have. On the other side, implicit type systems contain no explicit terms and thus no type annotations and no explicit reduction. Type systems always have an implicit version since it can be deduced from the explicit version by erasing type annotations. As a consequence, we assume the definition of the explicit version by default. Terms will refer to explicit terms, reduction to explicit reduction, and so on. In case of ambiguity, we will specify whether we are referring to the implicit or explicit version.

In this chapter, we give as much as possible the explicit and implicit version of the considered type systems. For System F_η , there is no explicit version in the literature, however we define one very naturally. For **Constraint ML** the case is worse, an explicit version would be very hard to write since the typing rules are not local: the consistency of constraints is checked at top level. An explicit version of **Constraint ML** would need to reduce the consistency proof of the top-level set of constraints.

When defining type systems, we factorize both versions by writing the judgments in the following form $E \Rightarrow J$ where E is an explicit element witnessing the judgment J . The witness E is a syntactic representation of the derivation of J . For instance, the explicit term judgment is of the form $M \Rightarrow \Gamma \vdash a : \tau$ where M is an explicit term which is a partial proof that the term a has type τ under environment Γ . It is a partial proof because it only contains enough information to rebuild the term a , and enough information to rebuild the type τ given an environment Γ . Explicit judgments are equivalent to implicit judgments. For instance, $M \Rightarrow \Gamma \vdash a : \tau$ holds for some M if and only if $\Gamma \vdash a : \tau$ holds.

3.1 The STLC

The Simply Typed Lambda Calculus (abbreviated as STLC) is the most simple type system. It provides both soundness and termination and illustrates all there is to know about the λ -calculus.

3.1.1 Definition

The STLC defines syntactical types, described on Figure 3.1. Types are written τ , σ , or ρ . They contain arrows $\tau \rightarrow \tau$ and products $\tau \times \tau$. Arrow types are used to type functions, while product types are used to type pairs. A term of type $\tau \rightarrow \sigma$ is a function with a domain represented by τ and a range represented by σ . While a term of type $\tau \times \sigma$ is a pair of a term of type τ and a term of type σ . Base types (leaves for the syntax tree of types), like `int` or `bool`, are necessary to write any type at all, but since they are neither necessary nor interesting to the theory, we leave them aside and use meta-variables in a generic way.

In order to type functions we need the notion of environments, written Γ . They assign a type to each free term variable. They are lists of bindings of the form $(x : \tau)$. The order does not matter yet, but it will in later type systems.

$E ::= \lambda(x : \tau) \square \mid \square M \mid M \square \mid \langle \square, M \rangle \mid \langle M, \square \rangle \mid \text{fst } \square \mid \text{snd } \square$	Evaluation contexts
$p ::= x \mid p v \mid \text{fst } p \mid \text{snd } p$	Prevalues
$v ::= p \mid \lambda(x : \tau) v \mid \langle v, v \rangle$	Values

Figure 3.2: STLC notations

$\frac{\text{REDCTX} \quad M \rightsquigarrow_{\beta} N}{E[M] \rightsquigarrow_{\beta} E[N]}$	$\text{REDAPP} \quad (\lambda(x : \tau) M) N \rightsquigarrow_{\beta} M[x/N]$	$\text{REDFST} \quad \text{fst } \langle M, N \rangle \rightsquigarrow_{\beta} M$	$\text{REDSND} \quad \text{snd } \langle M, N \rangle \rightsquigarrow_{\beta} N$
---	---	---	---

Figure 3.3: STLC reduction relation

Terms of the STLC mainly correspond to the terms of the λ -calculus and are given on Figure 3.1. They are written M or N and contain variables x , abstractions $\lambda(x : \tau) M$ where the term variable x is annotated with its type τ , applications $M M$, pairs $\langle M, M \rangle$, and projections $\text{fst } M$ and $\text{snd } M$.

We also have to define some notations, which are given on Figure 3.2. These notations are about the reduction of terms. Since the λ -calculus uses strong reduction, we need to use strong reduction for explicit terms too. And thus we define evaluation contexts, prevalues, and values for explicit terms in strong reduction.

Evaluation contexts, using the overloaded notation E , resemble those of the λ -calculus modulo the type annotation for term abstraction. They are one-hole contexts of depth one. Prevalues, again using the overloaded notation p , are variables or destructors applied to values but where a constructor is expected, in which case only a prevalue is applied. And values, using the overloaded notation v , are prevalues or constructors applied to values.

We define the reduction rules of explicit terms. These are given on Figure 3.3. We write $M \rightsquigarrow_{\beta} N$ to say that M reduces on N . Reduction rules resemble those of the λ -calculus modulo type annotations, which are ignored. Rule REDCTX is the context rule of strong reduction. Rule REDAPP tells that the application of a term N to an abstraction $\lambda(x : \tau) M$ reduces to the substitution of x by N in M , which we write $M[x/N]$. Finally, rules REDFST and REDSND tell that first and second projections of a pair reduce to the first and second elements of the pair respectively.

The implicit judgment $\Gamma \vdash a : \tau$ tells that the term a has type τ under environment Γ . The rules to derive this judgment are given on Figure 3.4. In the explicit version we add the term M witnessing the derivation: $M \Rightarrow \Gamma \vdash a : \tau$. This notation is actually coherent with our more general notion of explicit judgments written $E \Rightarrow J$ where E is a witness of the derivation of the judgment J . We reuse this notation later for coercions $G \Rightarrow \Gamma \vdash \tau \triangleright \sigma$ where the coercion G witnesses the coercion judgment $\Gamma \vdash \tau \triangleright \sigma$.

Rule TERMVAR tells that the term variable x has type τ under the well-formed environment Γ if the variable x is associated to the type τ in the environment Γ . The term witnessing this proof is the term variable x itself. Rule TERMLAM gives the function $\lambda x a$ the type $\tau \rightarrow \sigma$ under environment Γ , if the term a has type σ under the extended environment $\Gamma, (x : \tau)$. The environment is extended because x is now free in a and we need to say how it behaves. If we name M the term for $\Gamma, (x : \tau) \vdash a : \sigma$, then the term we use for the conclusion is $\lambda(x : \tau) M$ since the type τ cannot be guessed by looking at the term M only. Rule TERMAPP

$$\begin{array}{c}
\text{TERMVAR} \\
\frac{\Gamma \text{ env} \quad (x : \tau) \in \Gamma}{x \Rightarrow \Gamma \vdash x : \tau} \\
\\
\text{TERMLAM} \\
\frac{M \Rightarrow \Gamma, (x : \tau) \vdash a : \sigma}{\lambda(x : \tau) M \Rightarrow \Gamma \vdash \lambda x a : \tau \rightarrow \sigma} \\
\\
\text{TERMAPP} \\
\frac{M \Rightarrow \Gamma \vdash a : \tau \rightarrow \sigma \quad N \Rightarrow \Gamma \vdash b : \tau}{MN \Rightarrow \Gamma \vdash ab : \sigma} \\
\\
\text{TERMPAIR} \\
\frac{M \Rightarrow \Gamma \vdash a : \tau \quad N \Rightarrow \Gamma \vdash b : \sigma}{\langle M, N \rangle \Rightarrow \Gamma \vdash \langle a, b \rangle : \tau \times \sigma} \\
\\
\text{TERMFST} \\
\frac{M \Rightarrow \Gamma \vdash a : \tau \times \sigma}{\text{fst } M \Rightarrow \Gamma \vdash \text{fst } a : \tau} \\
\\
\text{TERMSND} \\
\frac{M \Rightarrow \Gamma \vdash a : \tau \times \sigma}{\text{snd } M \Rightarrow \Gamma \vdash \text{snd } a : \sigma}
\end{array}$$

Figure 3.4: STLC term judgment relation

$$\begin{array}{c}
\text{TYPEARR} \\
\frac{\Gamma \vdash \tau \text{ type} \quad \Gamma \vdash \sigma \text{ type}}{\Gamma \vdash \tau \rightarrow \sigma \text{ type}} \\
\\
\text{TYPEPROD} \\
\frac{\Gamma \vdash \tau \text{ type} \quad \Gamma \vdash \sigma \text{ type}}{\Gamma \vdash \tau \times \sigma \text{ type}} \\
\\
\text{ENVEMPTY} \\
\frac{}{\emptyset \text{ env}} \\
\\
\text{ENVTERM} \\
\frac{x \notin \text{dom}(\Gamma) \quad \Gamma \vdash \tau \text{ type}}{\Gamma, (x : \tau) \text{ env}}
\end{array}$$

Figure 3.5: STLC well formedness relations

tells that if a is a function of type $\tau \rightarrow \sigma$ under environment Γ and b is a term of type τ under environment Γ , then the application of b to a is of type σ under Γ . The environments are unchanged since the free variables of ab contain those of a and b . The term we use is MN , given the terms for the premises were M for the function and N for the argument.

The next three rules are about pairs and the environment does not play a role since no abstractions take place. Rule `TERMPAIR` tells that the pair $\langle a, b \rangle$ has type $\tau \times \sigma$ if a has type τ and b has type σ . The term is similar to the lambda term. Rule `TERMFST` (resp. `TERMSND`) tells that the first (resp. second) projection $\text{fst } a$ (resp. $\text{snd } a$) of a has type τ (resp. σ) if a has type $\tau \times \sigma$. The term is similar to the lambda term.

The judgment $\Gamma \vdash \tau \text{ type}$ tells that the type τ is well-formed under environment Γ . And judgment $\Gamma \text{ env}$ tells that the environment Γ is well-formed. The rules to derive these judgments are given on Figure 3.5. A type is always well-formed and an environment is well-formed if it binds its variable at most once and their types are well-formed.

Rule `TYPEARR` (resp. `TYPEPROD`) tells that $\tau \rightarrow \sigma$ (resp. $\tau \times \sigma$) is well-formed under environment Γ if both types τ and σ are well-formed under Γ . The empty environment \emptyset is well-formed by `ENVEMPTY` and the extension of the environment Γ with the binding $(x : \tau)$ is well-formed, by rule `ENVTERM`, if x is not already bound in Γ and τ is well-formed under Γ .

$$\begin{array}{ll}
[x] = x & \llbracket \langle M, N \rangle \rrbracket = \langle \llbracket M \rrbracket, \llbracket N \rrbracket \rangle \\
\llbracket \lambda(x : \tau) M \rrbracket = \lambda x \llbracket M \rrbracket & \llbracket \text{fst } M \rrbracket = \text{fst } \llbracket M \rrbracket \\
\llbracket M N \rrbracket = \llbracket M \rrbracket \llbracket N \rrbracket & \llbracket \text{snd } M \rrbracket = \text{snd } \llbracket M \rrbracket
\end{array}$$

Figure 3.6: STLC drop function

3.1.2 Properties

In an explicit type system, a judgment has to be unique according to its explicit entity. For instance, when the explicit term judgment $M \Rightarrow \Gamma \vdash a : \tau$ holds, we know that a is determined by M only and τ is a function of the term M and the environment Γ .

Lemma 4 (Uniqueness). *The following assertions hold.*

- If $M \Rightarrow \Gamma_1 \vdash a_1 : \tau_1$ and $M \Rightarrow \Gamma_2 \vdash a_2 : \tau_2$ hold, then $a_1 = a_2$ holds.
- If $M \Rightarrow \Gamma \vdash a : \tau_1$ and $M \Rightarrow \Gamma \vdash a : \tau_2$ hold, then $\tau_1 = \tau_2$ hold.

Actually a is a function of M even if M is not well-typed. We call this function the *drop* function and we write it $\llbracket M \rrbracket$. It is simply defined by dropping all type decorations in M . The formal definition is given on Figure 3.6. This lemma explains why we usually omit a in explicit term judgments.

Lemma 5. *If $M \Rightarrow \Gamma \vdash a : \tau$ holds, then $a = \llbracket M \rrbracket$ holds.*

Although we can write $\Gamma \vdash M : \tau$ instead of $M \Rightarrow \Gamma \vdash a : \tau$, we will not do so in order to keep similar notations between all type systems. In particular, some type systems have other explicit judgments than the term judgment and the notation $E \Rightarrow J$ where E is an explicit entity and J is an implicit judgment will be used for them.

The next lemma gives the equivalence between the implicit and explicit version of the type system. It tells that M is actually a function of the implicit typing derivation. In other words, we can extract from a derivation of $\Gamma \vdash a : \tau$ the term M such that $M \Rightarrow \Gamma \vdash a : \tau$ holds. And reciprocally that the term M is only a decoration and the validity of the judgment does not depend on it.

Lemma 6 (Equivalence). *$\Gamma \vdash a : \tau$ holds if and only if $M \Rightarrow \Gamma \vdash a : \tau$ holds for some M .*

We have some simulation properties about the explicit term reduction. If an explicit term can reduce, then its dropped lambda term can do the same reduction. And reciprocally, if the dropped lambda term of an explicit term can reduce, then the explicit term can do it too.

Lemma 7 (Bisimulation). *The following assertions hold.*

- If $M \rightsquigarrow_\beta N$ holds, then $\llbracket M \rrbracket \rightsquigarrow \llbracket N \rrbracket$ holds too.
- If $\llbracket M \rrbracket \rightsquigarrow b$ holds, then there is an N such that $M \rightsquigarrow_\beta N$ and $b = \llbracket N \rrbracket$ hold.
- If M is a value, then $\llbracket M \rrbracket$ is a value too.

Contrary to the untyped reduction relation of the λ -calculus, the explicit reduction relation is strongly normalizing for well-typed terms. And by bisimulation, well-typed lambda terms strongly normalize too.

Lemma 8 (Termination). *If $M \Rightarrow \Gamma \vdash a : \tau$ holds, then M strongly normalizes.*

Corollary 9. *If $\Gamma \vdash a : \tau$ holds, then a strongly normalizes.*

Similarly to the reduction relation of the λ -calculus, the explicit reduction relation is confluent. If an explicit term can reduce in two manners, then there is a term that joins these two reduction paths.

Lemma 10 (Confluence). *The explicit reduction is confluent.*

The STLC type system obeys some usual syntactical properties. The first property is called weakening. It tells that a well-formed type (resp. a well-typed term) under the environment Γ is also well-formed (resp. well-typed with the same type) under a well-formed extended environment Γ' . This property is in part used to prove the substitution lemma.

Lemma 11 (Weakening). *If $\Gamma \subseteq \Gamma'$ and Γ' env hold, then the following assertions hold:*

- *If $\Gamma \vdash \tau$ type holds, then $\Gamma' \vdash \tau$ type holds.*
- *If $M \Rightarrow \Gamma \vdash a : \tau$ holds, then $M \Rightarrow \Gamma' \vdash a : \tau$ holds.*

Corollary 12. *If Γ env and $(x : \tau) \in \Gamma$ hold, then $\Gamma \vdash \tau$ type holds.*

The second property is called substitution. It tells that the substitution $a[x/b]$ of x by b in a has type σ under environment Γ extended with Γ' if the argument b has type τ under Γ and the body a has type σ under Γ extended with x associated to τ and Γ' . This lemma is used to prove the subject reduction property for the REDAPP case.

Lemma 13 (Substitution). *If $M \Rightarrow \Gamma, (x : \tau), \Gamma' \vdash a : \sigma$ and $N \Rightarrow \Gamma \vdash b : \tau$ hold, then $M[x/N] \Rightarrow \Gamma, \Gamma' \vdash a[x/b] : \sigma$ holds.*

The next property is called extraction. It tells that the subcomponents of a judgment are well-formed. We can extract the well-formedness of the type and environment of a term judgment, and we can extract the well-formedness of the environment of a type judgment.

Lemma 14 (Extraction). *The following assertions hold:*

- *If $\Gamma \vdash \tau$ type holds, then Γ env holds.*
- *If $\Gamma \vdash a : \tau$ holds, then $\Gamma \vdash \tau$ type holds.*

The explicit reduction preserves typing. When a term a has type τ under environment Γ with witness M which reduces to N , then the term N actually witnesses that $\lfloor N \rfloor$ has type τ under Γ . By bisimulation we deduce the subject reduction property for the well-typed lambda terms. If the term a is well-typed and reduces to b , then it is still well-typed with the same typing (environment and type).

Lemma 15 (Explicit subject reduction). *If $M \Rightarrow \Gamma \vdash a : \tau$ and $M \rightsquigarrow_{\beta} N$ hold, then $N \Rightarrow \Gamma \vdash b : \tau$ holds where $b = \lfloor N \rfloor$.*

Corollary 16 (Subject reduction). *If $\Gamma \vdash a : \tau$ and $a \rightsquigarrow b$ hold, then $\Gamma \vdash b : \tau$ holds.*

The subject reduction property usually goes with the progress property. Both properties together show type soundness. Progress tells that if a term a is well-typed then it either reduces or it is a value. The same property holds for the explicit language. Progress is usually shown for closed terms. Since we are in a strong reduction setting, we state the progress lemma for all terms, including open terms. The result for closed term is a corollary. The proof of progress relies on a classification lemma, that gives the form of the typing derivation of values according to their type.

Lemma 17 (Explicit progress). *If $M \Rightarrow \Gamma \vdash a : \tau$ holds and M is not a value, then there is an term N such that $M \rightsquigarrow_{\beta} N$ holds.*

Lemma 18 (Progress). *If $\Gamma \vdash a : \tau$ holds and a is not a value, then there is a term b such that $a \rightsquigarrow b$ holds.*

And finally, the STLC type system obeys the soundness property: well-typed terms are sound. A term is sound if all its reductions do not encounter an error. This proposition relies on subject reduction and the fact that well-typed term are valid terms.

Proposition 19 (Soundness). *If $\Gamma \vdash a : \tau$ holds, then a is sound.*

3.2 System F

The STLC type system is sound but quite limited. For instance, the encodings we used in Section 2.4 cannot be given a generic type in the STLC, but they admit one in System F. A more simple example is the following sound and terminating term, which is typable in System F but not in the STLC:

$$\text{let } x = \lambda y y \text{ in } x x$$

According to the body of the `let`-construct, the type of x has to be an arrow type and the left-hand side of the arrow type has to be again the type of x . As a consequence we would need some sort of infinite type. One solution is to use recursive types as in System F_{rec} (see Section 3.3), while another one is to use polymorphism as described in this section.

Types are used to show particular properties (like soundness, termination, etc.) but they can also be used by the programmer to describe the behavior of a term, as documentation. For instance a term of type $\tau \rightarrow \sigma$ takes any term of type τ and returns a term of type σ . The type syntax of the STLC is quite poor to be used as documentation, we might want to extend it. A useful and powerful extension is polymorphism. The STLC extended with polymorphism is called System F.

Let's consider the identity term written `id` which definition is $\lambda x x$. The most we can say is that it takes a term and returns it as is. In the STLC we can say that for some type τ , if it takes a term of type τ , it returns a term of type τ . However it is true for any type τ but we cannot say it in the syntax of types. We need to introduce a new type constructor that does not have a computational content but only a typing (and documentation) content.

3.2.1 Definition

System F extends the syntax of the STLC with a class of type variables, written α or β . Syntactical types are also extended with type variables α and polymorphic types $\forall \alpha \tau$. Bindings

α, β	Type variables
$M, N ::= x \mid \lambda(x : \tau) M \mid M M \mid \langle M, M \rangle \mid \mathbf{fst} M \mid \mathbf{snd} M$ $\mid \Lambda \alpha M \mid M[\tau]$	Explicit terms
$\tau, \sigma, \rho ::= \alpha \mid \tau \rightarrow \tau \mid \tau \times \tau \mid \forall \alpha \tau$	Types
$\Gamma ::= \emptyset \mid \Gamma, (x : \tau) \mid \Gamma, \alpha$	Environments

Figure 3.7: System F syntax

$E ::= \lambda(x : \tau) [] \mid [] M \mid M [] \mid \langle [], M \rangle \mid \langle M, [] \rangle \mid \mathbf{fst} [] \mid \mathbf{snd} []$ $\mid \Lambda \alpha [] \mid [][\tau]$	Evaluation contexts
$p ::= x \mid p v \mid \mathbf{fst} p \mid \mathbf{snd} p \mid p[\tau]$	Prevalues
$v ::= p \mid \lambda(x : \tau) v \mid \langle v, v \rangle \mid \Lambda \alpha v$	Values

Figure 3.8: System F notations

are extended with type bindings α . Terms are extended with type abstractions $\Lambda \alpha M$ and type applications $M[\tau]$. All these changes are formally given on Figure 3.7.

Notations for System F follow the same logic as the λ -calculus or the STLC. They are given on Figure 3.8. Evaluation contexts are all one-hole terms of depth one and correspond to strong reduction. Prevalues are variables or destructors applied to values but where a constructor is expected, in which case only a prevalue is applied. And values are prevalues or constructors applied to values. The new prevalue is $p[\tau]$ since type application is a destructor and the new value is $\Lambda \alpha v$ since type abstraction is a constructor.

Reduction rules for System F are given on Figure 3.9. We write $M \rightsquigarrow_{\beta} N$ when M reduces to N with a computational content and $M \rightsquigarrow_{\iota} N$ when M reduces on N without changing the computational content of M . We use the meta-variable $\beta\iota$ to stand for β or ι . This labeling is used in the bisimulation property when linking the explicit reduction and the λ -calculus reduction.

The first four rules are those of the STLC. The last rule, namely REDFOR, is new and about polymorphism. Hence it is a ι -reduction rule while the STLC reduction rules were β -reduction rules. When a type application follows a type abstraction, reduction substitutes the term to which they are applied replacing the type variable with the type argument.

The term judgment relation is defined on Figure 3.10. The first six rules are those of the STLC. The last two rules are about polymorphism. Rule TERMG_{EN} tells that if a has type τ under the environment Γ extended with the type variable α , then it can also be typed $\forall \alpha \tau$ under Γ . The witness used for this typing rule is $\Lambda \alpha M$ when M is the witness for the unique premise. Rule TERM_{INST} does the contrary, it instantiates a type variable. If a has polymorphic type $\forall \alpha \tau$ and σ is a well-formed type, then a also has type τ where all free occurrences of α are substituted with σ , namely $\tau[\alpha/\sigma]$. The witness of the rule is $M[\sigma]$ where M is the witness of the unique term judgment premise.

Finally the well-formedness relations are given on Figure 3.11. Rule TYPEVAR tells that a type variable is well-formed if it is bound in the environment, which has to be well-formed. Rule TYPEARR (resp. TYPEPROD) tells that $\tau \rightarrow \sigma$ (resp. $\tau \times \sigma$) is well-formed if both τ and σ are well-formed. The polymorphic type $\forall \alpha \tau$ is well-formed if its body τ is well-formed in the same environment extended with α , according to rule TYPEFOR.

$$\begin{array}{c}
\text{REDCTX} \\
\frac{M \rightsquigarrow_{\beta\iota} N}{E[M] \rightsquigarrow_{\beta\iota} E[N]}
\end{array}
\quad
\begin{array}{c}
\text{REDAPP} \\
(\lambda(x : \tau) M) N \rightsquigarrow_{\beta} M[x/N]
\end{array}
\quad
\begin{array}{c}
\text{REDFST} \\
\text{fst } \langle M, N \rangle \rightsquigarrow_{\beta} M
\end{array}
\quad
\begin{array}{c}
\text{REDSND} \\
\text{snd } \langle M, N \rangle \rightsquigarrow_{\beta} N
\end{array}$$

$$\begin{array}{c}
\text{REDFOR} \\
(\Lambda\alpha M)[\tau] \rightsquigarrow_{\iota} M[\alpha/\tau]
\end{array}$$

Figure 3.9: System F reduction rules

$$\begin{array}{c}
\text{TERMVAR} \\
\frac{\Gamma \text{ env} \quad (x : \tau) \in \Gamma}{x \Rightarrow \Gamma \vdash x : \tau}
\end{array}
\quad
\begin{array}{c}
\text{TERMLAM} \\
\frac{M \Rightarrow \Gamma, (x : \tau) \vdash a : \sigma}{\lambda(x : \tau) M \Rightarrow \Gamma \vdash \lambda x a : \tau \rightarrow \sigma}
\end{array}$$

$$\begin{array}{c}
\text{TERMAPP} \\
\frac{M \Rightarrow \Gamma \vdash a : \tau \rightarrow \sigma \quad N \Rightarrow \Gamma \vdash b : \tau}{MN \Rightarrow \Gamma \vdash ab : \sigma}
\end{array}
\quad
\begin{array}{c}
\text{TERMPAIR} \\
\frac{M \Rightarrow \Gamma \vdash a : \tau \quad N \Rightarrow \Gamma \vdash b : \sigma}{\langle M, N \rangle \Rightarrow \Gamma \vdash \langle a, b \rangle : \tau \times \sigma}
\end{array}$$

$$\begin{array}{c}
\text{TERMFST} \\
\frac{M \Rightarrow \Gamma \vdash a : \tau \times \sigma}{\text{fst } M \Rightarrow \Gamma \vdash \text{fst } a : \tau}
\end{array}
\quad
\begin{array}{c}
\text{TERMSND} \\
\frac{M \Rightarrow \Gamma \vdash a : \tau \times \sigma}{\text{snd } M \Rightarrow \Gamma \vdash \text{snd } a : \sigma}
\end{array}$$

$$\begin{array}{c}
\text{TERMGEN} \\
\frac{M \Rightarrow \Gamma, \alpha \vdash a : \tau}{\Lambda\alpha M \Rightarrow \Gamma \vdash a : \forall\alpha \tau}
\end{array}
\quad
\begin{array}{c}
\text{TERMINST} \\
\frac{M \Rightarrow \Gamma \vdash a : \forall\alpha \tau \quad \Gamma \vdash \sigma \text{ type}}{M[\sigma] \Rightarrow \Gamma \vdash a : \tau[\alpha/\sigma]}
\end{array}$$

Figure 3.10: System F term judgment relation

The empty environment is always well-formed, rule `ENVEMPTY`. An environment extended with a term binding is well-formed, by rule `ENVTERM`, if its term variable is not already bound and its type is well-formed. A well-formed environment extended with a type binding is well-formed, by rule `ENVTYPE`, if its type variable is not already bound.

3.2.2 Properties

System F obeys almost the same properties as the STLC. The explicit term assures the uniqueness of the implicit term and its type when the term is well-typed. Concretely, if $M \Rightarrow \Gamma \vdash a : \tau$ holds, then a is a function of M only and τ is a function of both M and Γ .

Lemma 20 (Uniqueness). *The following assertions hold.*

- If $M \Rightarrow \Gamma_1 \vdash a_1 : \tau_1$ and $M \Rightarrow \Gamma_2 \vdash a_2 : \tau_2$ hold, then $a_1 = a_2$ holds.
- If $M \Rightarrow \Gamma \vdash a : \tau_1$ and $M \Rightarrow \Gamma \vdash a : \tau_2$ hold, then $\tau_1 = \tau_2$ hold.

Similarly to the STLC, a is a function of M even if M is not well-typed. We define the drop function of System F on Figure 3.12. The drop function drops the typing annotations. Type abstractions and type applications are erased. This lemma explain why we usually omit a in explicit term judgments.

$$\begin{array}{c}
\text{TYPEVAR} \\
\frac{\Gamma \text{ env} \quad \alpha \in \Gamma}{\Gamma \vdash \alpha \text{ type}} \\
\\
\text{ENVEMPTY} \\
\frac{}{\emptyset \text{ env}} \\
\\
\text{TYPEARR} \\
\frac{\Gamma \vdash \tau \text{ type} \quad \Gamma \vdash \sigma \text{ type}}{\Gamma \vdash \tau \rightarrow \sigma \text{ type}} \\
\\
\text{ENVTERM} \\
\frac{x \notin \text{dom}(\Gamma) \quad \Gamma \vdash \tau \text{ type}}{\Gamma, (x : \tau) \text{ env}} \\
\\
\text{TYPEPROD} \\
\frac{\Gamma \vdash \tau \text{ type} \quad \Gamma \vdash \sigma \text{ type}}{\Gamma \vdash \tau \times \sigma \text{ type}} \\
\\
\text{ENVTYPE} \\
\frac{\alpha \notin \text{dom}(\Gamma) \quad \Gamma \text{ env}}{\Gamma, \alpha \text{ env}} \\
\\
\text{TYPEFOR} \\
\frac{\Gamma, \alpha \vdash \tau \text{ type}}{\Gamma \vdash \forall \alpha \tau \text{ type}}
\end{array}$$

Figure 3.11: System F well-formedness relations

$$\begin{array}{ccc}
[x] = x & \langle [M, N] \rangle = \langle [M], [N] \rangle & [\Lambda \alpha M] = [M] \\
[\lambda(x : \tau) M] = \lambda x [M] & [\text{fst } M] = \text{fst } [M] & [M[\tau]] = [M] \\
[MN] = [M] [N] & [\text{snd } M] = \text{snd } [M] &
\end{array}$$

Figure 3.12: System F drop function

Lemma 21. *If $M \Rightarrow \Gamma \vdash a : \tau$ holds, then $a = [M]$ holds.*

System F also comes with an equivalence lemma as the STLC. The term M is a function of the implicit typing derivation. And the judgment holds also in absence of the explicit term.

Lemma 22 (Equivalence). *$\Gamma \vdash a : \tau$ holds if and only if $M \Rightarrow \Gamma \vdash a : \tau$ holds for some M .*

Lemma 23 (ι -termination). *The ι -reduction terminates.*

Proof. The number of explicit term nodes strictly decreases at each ι -reduction step. Only the types get bigger. \square

The bisimulation property is now more interesting. If a term can reduce with a β -step, then its dropped lambda term can do the same reduction. If a term can reduce with a ι -step, then its dropped term remains the same. Reciprocally, if a dropped lambda term can reduce, the term can do a series of ι -steps to do the same β -step. However, the term now has to be well-typed since we need a classification lemma about ι -normal forms.

Lemma 24 (Bisimulation). *The following assertions hold.*

- *If $M \rightsquigarrow_{\beta} N$ holds, then $[M] \rightsquigarrow [N]$ holds too.*
- *If $M \rightsquigarrow_{\iota} N$ holds, then $[M] = [N]$ holds.*
- *If $M \Rightarrow \Gamma \vdash a : \tau$ and $a \rightsquigarrow b$ holds, then there is an N such that $M \rightsquigarrow_{\iota}^* \rightsquigarrow_{\beta} N$ and $b = [N]$ hold.*
- *If M is an explicit value (see Figure 3.8), then $[M]$ is an implicit value.*

System F is strongly normalizing. The proof of this result is quite involved and initially due Girard [16]. Our proof of soundness for System F_{cc} in Chapter 5 is actually based on the ideas of this proof. The proof is semantic and interprets types as sets of strongly normalizing terms. If a term is well-typed, it is in its semantic type, and thus strongly normalizes.

Lemma 25 (Termination). *If $\Gamma \vdash a : \tau$ holds, then a strongly normalizes.*

Corollary 26. *If $M \Rightarrow \Gamma \vdash a : \tau$ holds, then M strongly normalizes.*

The explicit reduction of System F is also confluent. The proof can be done using strong normalization and local confluence, which holds since there are no critical pairs.

Lemma 27 (Confluence). *The explicit reduction is confluent.*

System F also has a weakening lemma. It tells that a well-formed type (resp. a well-typed term) under Γ is also well-formed (resp. well-typed with the same type) under an extended environment Γ' .

Lemma 28 (Weakening). *If $\Gamma \subseteq \Gamma'$ and Γ' env hold, then the following assertions hold.*

- *If $\Gamma \vdash \tau$ type holds, then $\Gamma' \vdash \tau$ type holds.*
- *If $M \Rightarrow \Gamma \vdash a : \tau$ holds, then $M \Rightarrow \Gamma' \vdash a : \tau$ holds.*

The term substitution lemma tells that the substitution $a[x/b]$ has type σ under Γ if the argument b has type τ under Γ and the body a has type σ under Γ extended with x associated to τ .

Lemma 29 (Term substitution). *If $N \Rightarrow \Gamma \vdash b : \tau$ and $M \Rightarrow \Gamma, (x : \tau), \Gamma' \vdash a : \sigma$ hold, then $M[x/N] \Rightarrow \Gamma, \Gamma' \vdash a[x/b] : \sigma$ holds.*

There is a similar property for types, called type substitution. If a type σ is well formed under environment Γ , then well-typed terms (resp. well-formed types) under Γ, α, Γ' are well-typed (resp. well-formed) under $\Gamma, \Gamma'[\alpha/\sigma]$ after type substitution replacing occurrences of α by type σ . Notice that we have to substitute in the remaining environment too as type variable α may appear there too.

Lemma 30 (Type substitution). *If $\Gamma \vdash \sigma$ type holds, then the following assertions hold.*

- *If $\Gamma, \alpha, \Gamma' \vdash \tau$ type holds, then $\Gamma, \Gamma'[\alpha/\sigma] \vdash \tau[\alpha/\sigma]$ type holds.*
- *If $M \Rightarrow \Gamma, \alpha, \Gamma' \vdash a : \tau$ holds, then $M[\alpha/\sigma] \Rightarrow \Gamma, \Gamma'[\alpha/\sigma] \vdash a : \tau[\alpha/\sigma]$ holds.*

The extraction property tells that the subcomponents of a judgment are well-formed. This is similar to the STLC. We can extract environment well-formedness from type well-formedness, and type well-formedness from term derivations.

Lemma 31 (Extraction). *The following assertions hold:*

- *If $\Gamma \vdash \tau$ type holds, then Γ env holds.*
- *If $\Gamma \vdash a : \tau$ holds, then $\Gamma \vdash \tau$ type holds.*

System F has also an explicit and implicit subject reduction properties which tell that if a term M has type τ under environment Γ and reduces on term N , then N also has type τ under Γ . Similarly for implicit terms when a reduces on b .

Lemma 32 (Explicit subject reduction). *If $M \Rightarrow \Gamma \vdash a : \tau$ and $M \rightsquigarrow_{\beta_L} N$ hold, then $N \Rightarrow \Gamma \vdash b : \tau$ holds where $b = \lfloor N \rfloor$.*

α, β	Type variables
$M, N ::= x \mid \lambda(x : \tau) M \mid M M \mid \langle M, M \rangle \mid \text{fst } M \mid \text{snd } M$ $\quad \mid \Lambda \alpha M \mid M[\tau] \mid \text{fold}_{\mu\alpha\tau} M \mid \text{unfold } M$	Explicit terms
$\tau, \sigma, \rho ::= \alpha \mid \tau \rightarrow \tau \mid \tau \times \tau \mid \forall \alpha \tau \mid \mu \alpha \tau$	Types
$\Gamma ::= \emptyset \mid \Gamma, (x : \tau) \mid \Gamma, \alpha$	Environments
$\text{rec} ::= \text{NE} \mid \text{WF}$	well-foundedness

Figure 3.13: System F_{rec} syntax

Lemma 33 (Subject reduction). *If $\Gamma \vdash a : \tau$ and $a \rightsquigarrow b$ hold, then $\Gamma \vdash b : \tau$ holds.*

This property usually goes with the progress property. Both properties together show type soundness. Progress tells that if a term a is well-typed then it either reduces or is a value. This property holds for both the explicit and implicit version.

Lemma 34 (Explicit progress). *If $M \Rightarrow \Gamma \vdash a : \tau$ holds and M is not a value, then there is a term N such that $M \rightsquigarrow N$ holds.*

Lemma 35 (Progress). *If $\Gamma \vdash a : \tau$ holds and a is not a value, then there is a term b such that $a \rightsquigarrow b$ holds.*

And finally, System F obeys the soundness property. We recall that a type system is sound if all its well-typed terms are sound. And a term is sound if it cannot reduce to an error as defined in Figure 2.3.

Proposition 36 (Soundness). *If $\Gamma \vdash a : \tau$ holds, then a is sound.*

3.3 System F_{rec}

System F can be extended with recursive types (Chapter 21 of [29]). For instance the type $\mu \alpha \tau \times \alpha$, which corresponds to the infinite type $\tau \times (\tau \times \dots)$, contains terms that behave as streams of type τ . However one precaution has to be taken when dealing with recursive types: all recursive types do not necessarily have a meaning. For instance $\mu \alpha \alpha$ has no structure, it is an ill-formed recursive definition. Recursive type $\mu \alpha \tau$ has a meaning only if the functor associating α to τ is well-founded, *i.e.* occurrences of α occur only under a computational type. Computational types, defined in the introduction, are related to the reduction and terms, while erasable types are related to typing. Computational types classify head normal forms and are thus those of the STLC.

3.3.1 Definition

System F_{rec} extends the syntactical types of System F with recursive types $\mu \alpha \tau$. The syntax of terms is also extended with foldings $\text{fold}_{\mu\alpha\tau} a$ and unfoldings $\text{unfold } a$. Finally, we add a syntactic class for well-foundedness: **NE** are for non-expansive functors, and **WF** for well-founded ones (see Figure 3.17). All these are formally given on Figure 3.13.

Notations for System F_{rec} follow the same logic as for System F . They are given on Figure 3.14. Evaluation contexts are all one-hole terms of depth one and correspond to strong

$ \begin{aligned} E ::= & \lambda(x : \tau) \square \mid \square M \mid M \square \mid \langle \square, M \rangle \mid \langle M, \square \rangle \mid \text{fst } \square \mid \text{snd } \square \\ & \mid \Lambda \alpha \square \mid \square[\tau] \mid \text{fold}_{\mu\alpha\tau} \square \mid \text{unfold } \square \\ p ::= & x \mid p v \mid \text{fst } p \mid \text{snd } p \mid p[\tau] \mid \text{unfold } p \\ v ::= & p \mid \lambda(x : \tau) v \mid \langle v, v \rangle \mid \Lambda \alpha v \mid \text{fold}_{\mu\alpha\tau} v \end{aligned} $	Evaluation contexts Prevalues Values
--	--

Figure 3.14: System F_{rec} notations

$ \begin{array}{c} \text{REDCTX} \\ \frac{M \rightsquigarrow_{\beta\iota} N}{E[M] \rightsquigarrow_{\beta\iota} E[N]} \end{array} $	$ \begin{array}{c} \text{REDAPP} \\ (\lambda(x : \tau) M) N \rightsquigarrow_{\beta} M[x/N] \end{array} $	$ \begin{array}{c} \text{REDFST} \\ \text{fst } \langle M, N \rangle \rightsquigarrow_{\beta} M \end{array} $	$ \begin{array}{c} \text{REDSND} \\ \text{snd } \langle M, N \rangle \rightsquigarrow_{\beta} N \end{array} $
$ \begin{array}{c} \text{REDFOR} \\ (\Lambda \alpha M)[\tau] \rightsquigarrow_{\iota} M[\alpha/\tau] \end{array} $		$ \begin{array}{c} \text{REDREC} \\ \text{unfold } (\text{fold}_{\mu\alpha\tau} M) \rightsquigarrow_{\iota} M \end{array} $	

Figure 3.15: System F_{rec} reduction rules

reduction. Prevalues are variables or destructors applied to values but where a constructor is expected, in which case only a prevalue is applied. And values are prevalues or constructors applied to values. The new prevalue is $\text{unfold } p$ since unfolding is a destructing operation, while $\text{fold}_{\mu\alpha\tau} v$ is the new value since it is a constructor (see the reduction rule REDREC).

Reduction rules for System F_{rec} are given on Figure 3.15. The first five rules are those of System F . The last rule, namely Rule REDREC , is new and about recursive types. When a term has its type folded and unfolded, it is like nothing happened. It is a ι -reduction rule since it does not change the computational content of its term.

The term judgment relation is defined on Figure 3.16. The first eight rules are those of System F . The last two rules are about recursive types. They allow to fold and unfold the definition of a recursive type $\mu\alpha\tau$. As for polymorphism these rules are implicit since they do not change the computational content of the term they are typing, but only its invariant. Rule TERMfold tells that a term a with type $\tau[\alpha/\mu\alpha\tau]$ has also type $\mu\alpha\tau$ by folding the recursive definition. Notice that the explicit version of the term needs a type annotation to know which recursive type it has to fold. Rule TERMunfold does the opposite: a term a of type $\mu\alpha\tau$ has also type $\tau[\alpha/\mu\alpha\tau]$.

Compared to System F , we need a judgment about the well-foundedness of functors, which we use to tell when a recursive type is well-formed. This judgment is written $\alpha \mapsto \tau : \text{rec}$ and means that the functor associating the type variable α to the type τ is non-expansive (resp. well-founded) if rec is NE (resp. WF). Intuitively, a well-founded functor uses its variable only under computational types, while a non-expansive functor may use its variable anywhere. The well-foundedness relation is given on Figure 3.17.

Rule RECvar tells that the identity functor is non-expansive. Rule RECcst tells that constant functors are well-founded. They do not use their variable, which means that they do not have a role in the recursive definition containing them. Rule RECsub tells that well-founded functors are also non-expansive: we can forget their well-foundedness. Rules RECarr and RECprod tell that functors, which head type constructors are arrows and products, are well-founded if their subfunctors are non-expansive. This comes from the fact that arrows and products are computational types. Finally, rules RECfor and RECmu tell that functors,

$$\begin{array}{c}
\text{TERMVAR} \\
\frac{\Gamma \text{ env} \quad (x : \tau) \in \Gamma}{x \Rightarrow \Gamma \vdash x : \tau} \\
\\
\text{TERMLAM} \\
\frac{M \Rightarrow \Gamma, (x : \tau) \vdash a : \sigma}{\lambda(x : \tau) M \Rightarrow \Gamma \vdash \lambda x a : \tau \rightarrow \sigma} \\
\\
\text{TERMAPP} \\
\frac{M \Rightarrow \Gamma \vdash a : \tau \rightarrow \sigma \quad N \Rightarrow \Gamma \vdash b : \tau}{MN \Rightarrow \Gamma \vdash ab : \sigma} \\
\text{TERMPAIR} \\
\frac{M \Rightarrow \Gamma \vdash a : \tau \quad N \Rightarrow \Gamma \vdash b : \sigma}{\langle M, N \rangle \Rightarrow \Gamma \vdash \langle a, b \rangle : \tau \times \sigma} \\
\\
\text{TERMFST} \\
\frac{M \Rightarrow \Gamma \vdash a : \tau \times \sigma}{\text{fst } M \Rightarrow \Gamma \vdash \text{fst } a : \tau} \\
\text{TERMSND} \\
\frac{M \Rightarrow \Gamma \vdash a : \tau \times \sigma}{\text{snd } M \Rightarrow \Gamma \vdash \text{snd } a : \sigma} \\
\text{TERMGEN} \\
\frac{M \Rightarrow \Gamma, \alpha \vdash a : \tau}{\Lambda \alpha M \Rightarrow \Gamma \vdash a : \forall \alpha \tau} \\
\\
\text{TERMINST} \\
\frac{M \Rightarrow \Gamma \vdash a : \forall \alpha \tau \quad \Gamma \vdash \sigma \text{ type}}{M[\sigma] \Rightarrow \Gamma \vdash a : \tau[\alpha/\sigma]} \\
\text{TERMFOLD} \\
\frac{M \Rightarrow \Gamma \vdash a : \tau[\alpha/\mu\alpha\tau]}{\text{fold}_{\mu\alpha\tau} M \Rightarrow \Gamma \vdash a : \mu\alpha\tau} \\
\\
\text{TERMUNFOLD} \\
\frac{M \Rightarrow \Gamma \vdash a : \mu\alpha\tau}{\text{unfold } M \Rightarrow \Gamma \vdash a : \tau[\alpha/\mu\alpha\tau]}
\end{array}$$

Figure 3.16: System F_{rec} term judgment relation

$$\begin{array}{c}
\text{RECVAR} \\
\alpha \mapsto \alpha : \text{NE} \\
\\
\text{RECCST} \\
\frac{\alpha \notin \text{fv}(\tau)}{\alpha \mapsto \tau : \text{WF}} \\
\\
\text{RECSUB} \\
\frac{\alpha \mapsto \tau : \text{WF}}{\alpha \mapsto \tau : \text{NE}} \\
\\
\text{RECArr} \\
\frac{\alpha \mapsto \tau : \text{NE} \quad \alpha \mapsto \sigma : \text{NE}}{\alpha \mapsto \tau \rightarrow \sigma : \text{WF}} \\
\\
\text{RECPROD} \\
\frac{\alpha \mapsto \tau : \text{NE} \quad \alpha \mapsto \sigma : \text{NE}}{\alpha \mapsto \tau \times \sigma : \text{WF}} \\
\\
\text{RECFor} \\
\frac{\alpha \mapsto \tau : \text{rec}}{\alpha \mapsto \forall \beta \tau : \text{rec}} \\
\\
\text{RECMu} \\
\frac{\beta \mapsto \tau : \text{WF} \quad \alpha \mapsto \tau : \text{rec}}{\alpha \mapsto \mu\beta\tau : \text{rec}}
\end{array}$$

Figure 3.17: System F_{rec} well-foundedness judgment relation

which head type constructors are polymorphic types or recursive types, are non-expansive (resp. well-founded) if their subfunctor is non-expansive (resp. well-founded). These type constructors are erasable and not computational. Additionally for the functor ending with a recursive type, an additional premise is necessary for the recursive type to be also well-founded.

Finally the well-formedness relations are given on Figure 3.18. Rule `TYPEVAR` tells that a type variable is well-formedness if it is bound in the environment, which has to be well-formed. Rule `TYPEARR` (resp. `TYPEPROD`) tells that $\tau \rightarrow \sigma$ (resp. $\tau \times \sigma$) is well-formed if both τ and σ are well-formed. The polymorphic type $\forall \alpha \tau$ is well-formed if τ is well-formed in the same environment extended with α , according to rule `TYPEFOR`. Rule `TYPEMU` tells that a recursive type is well-formed if its functor is well-founded and its body is well-formed under its environment extended with its type variable.

The empty environment is always well-formed, rule `ENVEMPTY`. An environment extended with a term binding is well-formed, by rule `ENVTERM`, if its term variable is not already

$$\begin{array}{c}
\text{TYPEVAR} \\
\frac{\Gamma \text{ env} \quad \alpha \in \Gamma}{\Gamma \vdash \alpha \text{ type}} \\
\\
\text{TYPEARR} \\
\frac{\Gamma \vdash \tau \text{ type} \quad \Gamma \vdash \sigma \text{ type}}{\Gamma \vdash \tau \rightarrow \sigma \text{ type}} \\
\\
\text{TYPEPROD} \\
\frac{\Gamma \vdash \tau \text{ type} \quad \Gamma \vdash \sigma \text{ type}}{\Gamma \vdash \tau \times \sigma \text{ type}} \\
\\
\text{TYPEFOR} \\
\frac{\Gamma, \alpha \vdash \tau \text{ type}}{\Gamma \vdash \forall \alpha \tau \text{ type}} \\
\\
\text{TYPEMU} \\
\frac{\Gamma, \alpha \vdash \tau \text{ type} \quad \alpha \mapsto \tau : \text{WF}}{\Gamma \vdash \mu \alpha \tau \text{ type}} \\
\\
\text{ENVEMPTY} \\
\emptyset \text{ env} \\
\\
\text{ENVTERM} \\
\frac{x \notin \text{dom}(\Gamma) \quad \Gamma \vdash \tau \text{ type}}{\Gamma, (x : \tau) \text{ env}} \\
\\
\text{ENVTYPE} \\
\frac{\alpha \notin \text{dom}(\Gamma) \quad \Gamma \text{ env}}{\Gamma, \alpha \text{ env}}
\end{array}$$

Figure 3.18: System F_{rec} well-formedness relations

bound and its type is well-formed. A well-formed environment extended with a type binding is well-formed, by rule ENVTYPE , if its type variable is not already bound.

3.3.2 Properties

System F_{rec} obeys almost the same properties as System F . The main difference is about termination. While every well-typed terms of System F strongly normalize, some well-typed terms of System F_{rec} may loop indefinitely. However these terms remain sound. One example is the omega term $\text{omega} \stackrel{\text{def}}{=} (\lambda y y y) (\lambda x x x)$. For any well-formed type τ under environment Γ , we define the term N as $\lambda(y : \mu \alpha (\alpha \rightarrow \tau)) (\text{unfold } y) y$ and the term M as $N (\text{fold}_{\mu \alpha (\alpha \rightarrow \tau)} N)$. We have $\llbracket M \rrbracket = \text{omega}$. We can easily see that the term M defines a derivation for $\Gamma \vdash \text{omega} : \tau$. We did not only prove that omega is well-typed, but that it accepts any well-formed type, in particular $\forall \alpha \alpha$, which we usually call the bottom type.

Despite the non-termination of some terms, the ι -reduction still strongly normalizes for all terms even for ill-typed ones. This property is useful to prove the backward simulation property, which can be used to prove subject reduction for the explicit version.

Lemma 37. *The ι -reduction strongly normalizes.*

Proof. The number of explicit term nodes strictly decreases during reduction. \square

It has been shown that with restricted forms of recursion [23], System F_{rec} may recover strong normalization. The most simple restriction is to consider positive recursion, where recursive type variables may only be used in covariant occurrences.

Besides the different restrictions on recursive types to recover strong normalization, there is another notion about recursive types which is subject to variants. System F_{rec} can be presented with equi-recursive types or iso-recursive types. The version of this section has iso-recursive types: recursive types are isomorphic to their unfolding. With equi-recursive types, recursive types are equal to their unfolding. This last notion is more general and subsumes the first notion. See Section 5.4.4 for a discussion about equi- and iso-recursive types.

α, β		Type variables
$M, N ::= x \mid \lambda(x : \tau) M \mid M M \mid \langle M, M \rangle \mid \text{fst } M \mid \text{snd } M$		Explicit terms
$\quad \mid \Lambda \alpha M \mid M[\tau] \mid G\langle M \rangle$		
$\tau, \sigma, \rho ::= \alpha \mid \tau \rightarrow \tau \mid \tau \times \tau \mid \forall \alpha \tau$		Types
$G ::= \diamond \mid G \circ G \mid \Lambda \alpha \mid \cdot \tau \mid G \xrightarrow{\tau} G \mid G \times G \mid \forall \alpha G \mid \text{dist}$		Explicit containments
$\Gamma ::= \emptyset \mid \Gamma, (x : \tau) \mid \Gamma, \alpha$		Environments

Figure 3.19: System F_η syntax

3.4 System F_η

As types are approximations of terms and some approximations are finer than others (fewer terms satisfy them), we may want to have an order between types related to the order between approximations. Giving a syntax in the type system for this ordering is called subtyping or type containment. This mechanism was studied along polymorphism in System F_η by Mitchell [25]. Since System F_η was described implicitly, we adapt the presentation for our explicit version. The main difference is for the distributivity rule which was $\forall \alpha (\tau \rightarrow \sigma) \triangleright (\forall \alpha \tau) \rightarrow \forall \alpha \sigma$ and is now $\forall \alpha (\tau \rightarrow \sigma) \triangleright \tau \rightarrow \forall \alpha \sigma$, given that α is not free in τ . The original rule is derivable in our setting. We also add product types.

We want our ordering to satisfy the fact that, if a type τ is smaller than σ under Γ with proof G , written $G \Rightarrow \Gamma \vdash \tau \triangleright \sigma$, then all terms of type τ under Γ also have type σ under Γ . In the explicit version the term M has to be annotated with the proof G , written $G\langle M \rangle$, when such retyping occurs.

Mitchell proved that the implicit version of this type system is equivalent to a version without the containment relation and the containment typing rule, but with an additional typing rule about η -expansion. This typing rule allows to type a function by actually typing its η -expansion. We can extend this idea to product types as well. In other words, a also has type $\tau \rightarrow \sigma$ under Γ , if its η -expansion $\lambda x a x$ has the same type $\tau \rightarrow \sigma$ under the same environment Γ . The typing rule follows:

$$\begin{array}{c}
\text{TERMETAARR} \\
\frac{\Gamma \vdash \lambda x a x : \tau \rightarrow \sigma \quad x \notin \text{fv}(a)}{\Gamma \vdash a : \tau \rightarrow \sigma}
\end{array}
\qquad
\begin{array}{c}
\text{TERMETAPROD} \\
\frac{\Gamma \vdash \langle \text{fst } a, \text{snd } a \rangle : \tau \times \sigma}{\Gamma \vdash a : \tau \times \sigma}
\end{array}$$

This version actually explains the name of System F_η . This rule implies that a term a has type τ under Γ in System F_η , if there is a term b of type τ under Γ in System F such that b η -reduces to a .

3.4.1 Definition

System F_η extends System F . The syntax of terms is extended with containment applications $G\langle M \rangle$. A new syntactical class for containment proofs is added. Containment proofs are written G and contain the reflexivity proof \diamond , transitivity proofs $G \circ G$, free generalizations $\Lambda \alpha$, type applications $\cdot \tau$, arrow subtypings $G \xrightarrow{\tau} G$, product subtypings $G \times G$, polymorphism congruences $\forall \alpha G$, and distributivity proofs dist . These changes with respect to System F are formally given on Figure 3.19.

$ \begin{aligned} E ::= & \lambda(x : \tau) \square \mid \square M \mid M \square \mid \langle \square, M \rangle \mid \langle M, \square \rangle \mid \mathbf{fst} \square \mid \mathbf{snd} \square \\ & \mid \Lambda \alpha \square \mid \square[\tau] \mid G(\square) \\ p ::= & x \mid p v \mid \mathbf{fst} p \mid \mathbf{snd} p \mid p[\tau] \mid (\forall \alpha G)\langle p \rangle \\ & \mid \mathbf{dist}\langle p \rangle \mid \mathbf{dist}\langle \Lambda \alpha p \rangle \mid (G \xrightarrow{\tau} G)\langle p \rangle \mid (G \times G)\langle p \rangle \\ v ::= & p \mid \lambda(x : \tau) v \mid \langle v, v \rangle \mid \Lambda \alpha v \end{aligned} $	Evaluation contexts Prevalues Values
---	--

Figure 3.20: System F_η notations

$ \begin{array}{c} \text{RED CTX} \\ \frac{M \rightsquigarrow_\beta N}{E[M] \rightsquigarrow_\beta E[N]} \end{array} $	$ \begin{array}{c} \text{RED APP} \\ (\lambda(x : \tau) M) N \rightsquigarrow_\beta M[x/N] \end{array} $	$ \begin{array}{c} \text{RED FST} \\ \mathbf{fst} \langle M, N \rangle \rightsquigarrow_\beta M \end{array} $	$ \begin{array}{c} \text{RED SND} \\ \mathbf{snd} \langle M, N \rangle \rightsquigarrow_\beta N \end{array} $
$ \begin{array}{c} \text{RED FOR} \\ (\Lambda \alpha M)[\tau] \rightsquigarrow_\iota M[\alpha/\tau] \end{array} $	$ \begin{array}{c} \text{RED REFL} \\ \Diamond \langle M \rangle \rightsquigarrow_\iota M \end{array} $	$ \begin{array}{c} \text{RED TRANS} \\ (G_2 \circ G_1)\langle M \rangle \rightsquigarrow_\iota G_2\langle G_1\langle M \rangle \rangle \end{array} $	
	$ \begin{array}{c} \text{RED TLAM} \\ \Lambda \alpha \langle M \rangle \rightsquigarrow_\iota \Lambda \alpha M \end{array} $	$ \begin{array}{c} \text{RED TAPP} \\ \cdot \tau \langle M \rangle \rightsquigarrow_\iota M[\tau] \end{array} $	
	$ \begin{array}{c} \text{RED ARR} \\ (G_1 \xrightarrow{\tau} G_2)\langle \lambda(x : \tau') M \rangle \rightsquigarrow_\iota \lambda(x : \tau) G_2\langle M[x/G_1\langle x \rangle] \rangle \end{array} $		
$ \begin{array}{c} \text{RED PROD} \\ (G_1 \times G_2)\langle \langle M_1, M_2 \rangle \rangle \rightsquigarrow_\iota \langle G_1\langle M_1 \rangle, G_2\langle M_2 \rangle \rangle \end{array} $	$ \begin{array}{c} \text{RED CONGR} \\ (\forall \alpha G)\langle \Lambda \alpha M \rangle \rightsquigarrow_\iota \Lambda \alpha G\langle M \rangle \end{array} $		
$ \begin{array}{c} \text{RED DIST ARR} \\ \mathbf{dist}\langle \Lambda \alpha \lambda(x : \tau) M \rangle \rightsquigarrow_\iota \lambda(x : \tau) \Lambda \alpha M \end{array} $	$ \begin{array}{c} \text{RED DIST PROD} \\ \mathbf{dist}\langle \Lambda \alpha \langle M_1, M_2 \rangle \rangle \rightsquigarrow_\iota \langle \Lambda \alpha M_1, \Lambda \alpha M_2 \rangle \end{array} $		

Figure 3.21: System F_η reduction rules

The syntax of evaluation contexts, prevalues, and values for System F_η follow the same schema as System F . They are given on Figure 3.20. Evaluation contexts are all one-hole terms of depth one and define a strong notion of reduction. Prevalues are variables or destructors applied to prevalues. And values are prevalues or constructors applied to values. There are only new prevalues since containments proofs are only destructors in this setting. The polymorphism congruence waits for a type abstraction to reduce, so $(\forall \alpha G)\langle p \rangle$ is a prevalue. The distributivity proof waits for a type abstraction followed by a computational constructor (like a term abstraction or a pair). So both $\mathbf{dist}\langle p \rangle$ and $\mathbf{dist}\langle \Lambda \alpha p \rangle$ may not reduce and are prevalues. Finally computational subtypings wait a computational constructor, so both $(G \xrightarrow{\tau} G)\langle p \rangle$ and $(G \times G)\langle p \rangle$ are prevalues.

Reduction rules for System F_η are given on Figure 3.21. The first five rules are those of System F . The next rules are new: one for each containment proof. The distributivity proof has two reduction rules since it has one reduction rule per computational type: arrow and product types in our setting. Rule REDREFL shows that the reflexivity proof does not modify a term and hence its typing. Rule REDTRANS shows that coercing the term M with $G_2 \circ G_1$ is like coercing M with G_1 and then coercing the result with G_2 . The typing modifications are done in sequence. Rules REDTLAM and REDTAPP simply unfold the definition of the

$$\begin{array}{c}
\text{TERMVAR} \\
\frac{\Gamma \text{ env} \quad (x : \tau) \in \Gamma}{x \Rightarrow \Gamma \vdash x : \tau} \\
\\
\text{TERMLAM} \\
\frac{M \Rightarrow \Gamma, (x : \tau) \vdash a : \sigma}{\lambda(x : \tau) M \Rightarrow \Gamma \vdash \lambda x a : \tau \rightarrow \sigma} \\
\\
\text{TERMAPP} \\
\frac{M \Rightarrow \Gamma \vdash a : \tau \rightarrow \sigma \quad N \Rightarrow \Gamma \vdash b : \tau}{MN \Rightarrow \Gamma \vdash ab : \sigma} \\
\\
\text{TERMPAIR} \\
\frac{M \Rightarrow \Gamma \vdash a : \tau \quad N \Rightarrow \Gamma \vdash b : \sigma}{\langle M, N \rangle \Rightarrow \Gamma \vdash \langle a, b \rangle : \tau \times \sigma} \\
\\
\text{TERMFST} \\
\frac{M \Rightarrow \Gamma \vdash a : \tau \times \sigma}{\text{fst } M \Rightarrow \Gamma \vdash \text{fst } a : \tau} \\
\\
\text{TERMSND} \\
\frac{M \Rightarrow \Gamma \vdash a : \tau \times \sigma}{\text{snd } M \Rightarrow \Gamma \vdash \text{snd } a : \sigma} \\
\\
\text{TERMGEN} \\
\frac{M \Rightarrow \Gamma, \alpha \vdash a : \tau}{\Lambda \alpha M \Rightarrow \Gamma \vdash a : \forall \alpha \tau} \\
\\
\text{TERMINST} \\
\frac{M \Rightarrow \Gamma \vdash a : \forall \alpha \tau \quad \Gamma \vdash \sigma \text{ type}}{M[\sigma] \Rightarrow \Gamma \vdash a : \tau[\alpha/\sigma]} \\
\\
\text{TERMCONT} \\
\frac{M \Rightarrow \Gamma \vdash a : \tau \quad G \Rightarrow \Gamma \vdash \tau \triangleright \sigma}{G\langle M \rangle \Rightarrow \Gamma \vdash a : \sigma}
\end{array}$$

Figure 3.22: System F_η term judgment relation

generalization and instantiation containment proofs.

Rule REDARR expects a function and applies G_2 to its body and G_1 to its arguments. In doing so the type annotation of the abstraction changes from τ' to τ , which is why we had to annotate the arrow subtyping coercion. Rule REDPROD expects a pair and applies G_1 to its first component and G_2 to its second one. Finally, rule REDCONGR applies its subcontainment proof to the body of the type abstraction it is applied to. Rule REDDISTARR swaps type and term abstractions. And rule REDDISTPROD swaps type abstraction and pairs.

The term judgment relation is defined on Figure 3.22. The first eight rules are those of System F. The new rule is the last one about containment, named TERMCONT. It tells that if a term a has type τ under Γ with witness M , then it also has type σ under the same environment Γ with witness $G\langle M \rangle$, given that G witnesses the containment proof that τ is smaller than σ under Γ .

This containment relation is actually a particularity of System F_η . Its judgment is written $G \Rightarrow \Gamma \vdash \tau \triangleright \sigma$ and means that G is a containment proof that the type τ is smaller than the type σ under the environment Γ . The rules are given on Figure 3.23. Notice that the extraction property for this judgment (Lemma 40) is quite particular, not only the well-formedness of Γ is an input, but also the well-formedness of $\Gamma \vdash \tau \text{ type}$ (which implies the well-formedness of Γ) is given as an input. This explains why premises about the well-formedness of τ are absent. This extraction property also explains how we can factorize the distributivity proofs for the arrow and product types in one proof: if the input type is an arrow then **dist** stands for the arrow distributivity rule, and if the input type is a product then **dist** stands for the product distributivity rule. This factorization permits to have fewer type annotations on coercions and thus on terms.

Rule CONTREFL defines \diamond as the containment proof that any type is contained in itself. Similarly rule CONTTRANS defines $G_2 \circ G_1$ as the proof for containment transitivity. If G_1 proves that τ_2 contains τ_1 and G_2 proves that τ_3 contains τ_2 , then $G_2 \circ G_1$ proves that τ_3 contains τ_1 by transitivity.

Rule CONTTLAM tells that we can add a forall quantifier on any well-formed type as long

$$\begin{array}{c}
\text{CONTREFL} \\
\Diamond \Rightarrow \Gamma \vdash \tau \triangleright \tau \\
\\
\text{CONTTRANS} \\
\frac{G_1 \Rightarrow \Gamma \vdash \tau_1 \triangleright \tau_2 \quad G_2 \Rightarrow \Gamma \vdash \tau_2 \triangleright \tau_3}{G_2 \circ G_1 \Rightarrow \Gamma \vdash \tau_1 \triangleright \tau_3} \\
\\
\text{CONTTLAM} \\
\frac{\alpha \notin \text{dom}(\Gamma)}{\Lambda \alpha \Rightarrow \Gamma \vdash \tau \triangleright \forall \alpha \tau} \\
\\
\text{CONTTAPP} \\
\frac{\Gamma \vdash \sigma \text{ type}}{\cdot \sigma \Rightarrow \Gamma \vdash \forall \alpha \tau \triangleright \tau[\alpha/\sigma]} \\
\\
\text{CONTARR} \\
\frac{\Gamma \vdash \tau \text{ type} \quad G_1 \Rightarrow \Gamma \vdash \tau \triangleright \tau' \quad G_2 \Rightarrow \Gamma \vdash \sigma' \triangleright \sigma}{G_1 \xrightarrow{\tau} G_2 \Rightarrow \Gamma \vdash \tau' \rightarrow \sigma' \triangleright \tau \rightarrow \sigma} \\
\\
\text{CONTPROD} \\
\frac{G_1 \Rightarrow \Gamma \vdash \tau' \triangleright \tau \quad G_2 \Rightarrow \Gamma \vdash \sigma' \triangleright \sigma}{G_1 \times G_2 \Rightarrow \Gamma \vdash \tau' \times \sigma' \triangleright \tau \times \sigma} \\
\\
\text{CONTCONGR} \\
\frac{G \Rightarrow \Gamma, \alpha \vdash \tau \triangleright \sigma}{\forall \alpha G \Rightarrow \Gamma \vdash \forall \alpha \tau \triangleright \forall \alpha \sigma} \\
\\
\text{CONTDISTARR} \\
\frac{\Gamma \vdash \tau \text{ type}}{\text{dist} \Rightarrow \Gamma \vdash \forall \alpha (\tau \rightarrow \sigma) \triangleright \tau \rightarrow \forall \alpha \sigma} \\
\\
\text{CONTDISTPROD} \\
\text{dist} \Rightarrow \Gamma \vdash \forall \alpha (\tau \times \sigma) \triangleright \forall \alpha \tau \times \forall \alpha \sigma
\end{array}$$

Figure 3.23: System F_η containment judgment relation

as the type variable was not already bound. Said otherwise, we can quantify over any fresh type variable. This rule resemble `TERMGEN` with the main difference that, in `TERMGEN`, the type variable may be bound prior to quantification, whereas the containment rule only quantifies over fresh type variables. This difference will actually be removed in type systems System F_t^p and System F_{cc} (see Part II). Rule `COERTLAM` in Figure 4.11 has no restriction on the generalized type variable α , and is thus equivalent to rule `TERMGEN`. This is a crucial point if we want to define type systems as coercions, which is the technique we use in order to unify all the type systems of this chapter. Rule `TERMGEN` is about an erasable type, namely the polymorphic type, and should thus be in the coercion judgment and not in the typing judgment for terms.

Rule `CONTTAPP` is the analog of rule `TERMINST`. It adds the instantiation typing rule in the containment relation and this time there is no loss of power. A polymorphic type is smaller than any of its instantiations as long as the argument type is well-formed.

Rule `CONTARR` is a standard rule in type systems with subtyping. It is the arrow congruence rule, which means that it tells when an arrow type is smaller than another arrow type. Type $\tau' \rightarrow \sigma'$ is smaller than $\tau \rightarrow \sigma$ under Γ if τ is smaller than τ' under Γ and σ' is smaller than σ under Γ . Notice the inversion of inclusion for the domain types, this comes from the contravariance of the domain of the arrow type. This rule needs an additional premise for well-formedness: type τ has to be well-formed under Γ . Notice that the containment proof for this rule needs a type annotation. This annotation is needed for reduction to rebuild the abstraction type annotation in rule `REDARR`. Rule `CONTPROD` is similar. Type $\tau' \times \sigma'$ is smaller than $\tau \times \sigma$ under Γ if τ' is smaller than τ under Γ and σ' is smaller than σ under Γ .

Rule `CONTCONGR` is a congruence rule for the polymorphic type. If G is a proof that τ is contained in σ under the extended environment Γ, α , then $\forall \alpha G$ is a proof that $\forall \alpha \tau$ is contained in $\forall \alpha \sigma$ under Γ . This rule is the only rule binding a variable for a subproof.

The distributivity rule `CONTDISTARR` is quite particular to System F_η . Most type systems with subtyping do not have this rule. It tells that we can permute type abstraction and term abstraction as long as the term abstraction type does not depend on the abstract type.

- If $G \Rightarrow \Gamma \vdash \tau \triangleright \sigma_1$ and $G \Rightarrow \Gamma \vdash \tau \triangleright \sigma_2$ hold, then $\sigma_1 = \sigma_2$ holds.

The equivalence lemma also has to be extended. However, this extension is much more natural, because it follows our schema about explicit entities witnessing judgments. From an implicit containment judgment, we can create a proof G witnessing this judgment.

Lemma 39 (Equivalence). *The following assertions hold.*

- $\Gamma \vdash a : \tau$ holds if and only if $M \Rightarrow \Gamma \vdash a : \tau$ holds for some M .
- $\Gamma \vdash \tau \triangleright \sigma$ holds if and only if $G \Rightarrow \Gamma \vdash \tau \triangleright \sigma$ holds for some G .

Finally, the extraction lemma is extended for containments. A containment derivation can transform the well-formedness of its left-hand type to the well-formedness of its right-hand type.

Lemma 40 (Extraction). *The following assertions hold:*

- If $\Gamma \vdash \tau$ type holds, then Γ env holds.
- If $M \Rightarrow \Gamma \vdash a : \tau$ holds, then $\Gamma \vdash \tau$ type holds.
- If $G \Rightarrow \Gamma \vdash \tau \triangleright \sigma$ and $\Gamma \vdash \tau$ type hold, then $\Gamma \vdash \sigma$ type holds.

3.5 MLF

An important feature for surface type systems, which are type systems for the programmer, is inference. Inference allows programmers to avoid writing all the typing derivations but only parts of them. Ideally and without any documentation consideration, we might want the programmer to write only the lambda term he wants to run. And at most we want him to write the explicit term. While previous type systems (but the STLC) have no complete inference, MLF which is more expressive than System F has a complete inference as long as the function parameters that are used polymorphically are annotated. However, we do not present MLF in its surface type system version[18], but in its kernel version[32], which is posterior.

3.5.1 Definition

MLF extends the STLC with instance-bounded polymorphism (which we may also refer to as lower bounded polymorphism) and bottom type. The polymorphic type of MLF abstracts over a type variable that is an instance of a lower bound type, instead of just abstracting over a type variable as it is the case in System F. We recover the polymorphic type $\forall \alpha \tau$ of System F by using an instance-bounded polymorphic type $\forall (\alpha \triangleleft \perp) \tau$ with a particular lower bound called bottom. All types are instance of the bottom type, and thus we can instantiate a lower bounded polymorphic type with bound bottom $\forall (\alpha \triangleleft \perp) \tau$ with any type σ as we can do in System F, because we have $\perp \triangleright \sigma$. So MLF is an extension of System F by its features and an extension of the STLC by its syntax.

MLF adds type and instance variables to the STLC. Type variables are written α or β , while instance variables are written c_α . All instance variables are linked to a unique type variable. Terms are extended with lower bounded abstraction $\Lambda(\alpha \triangleleft c : \tau) M$ and instance application $G\langle M \rangle$. Lower bounded abstraction abstracts over both the type variable α and

α, β	Type variables
c_α	Instance variables
$M, N ::= x \mid \lambda(x : \tau) M \mid M M \mid \langle M, M \rangle \mid \mathbf{fst} M \mid \mathbf{snd} M$ $\mid \Lambda(\alpha \triangleleft c : \tau) M \mid G\langle M \rangle$	Explicit terms
$\tau, \sigma, \rho ::= \alpha \mid \tau \rightarrow \tau \mid \tau \times \tau \mid \forall(\alpha \triangleleft \tau) \tau \mid \perp$	Types
$G ::= \perp\tau \mid c_\alpha \mid \forall(\triangleleft G) \mid \forall(\alpha \triangleleft c :) G \mid \& \mid \mathfrak{A} \mid G \circ G \mid \diamond$	Instantiations
$\Gamma ::= \emptyset \mid \Gamma, (x : \tau) \mid \Gamma, \alpha \mid \Gamma, (c_\alpha : \tau \triangleright \tau)$	Environments

Figure 3.25: MLF syntax

$E ::= \lambda(x : \tau) [] \mid [] M \mid M [] \mid \langle [], M \rangle \mid \langle M, [] \rangle \mid \mathbf{fst} [] \mid \mathbf{snd} []$ $\mid \Lambda(\alpha \triangleleft c : \tau) [] \mid G\langle [] \rangle$	Evaluation contexts
$p ::= x \mid p v \mid \mathbf{fst} p \mid \mathbf{snd} p \mid \&\langle p \rangle$ $\mid (\forall(\alpha \triangleleft c :) G)\langle p \rangle \mid (\forall(\triangleleft G))\langle p \rangle \mid \perp\tau\langle p \rangle \mid c_\alpha\langle v \rangle$	Prevalues
$v ::= p \mid \lambda(x : \tau) v \mid \langle v, v \rangle \mid \Lambda(\alpha \triangleleft c : \tau) v$	Values

Figure 3.26: MLF notations

the instance variable c of instance type $\tau \triangleright \alpha$. The instance variable c is associated to the type variable α , and we refer to it as c_α to enhance this association. The instance application uses an instance proof G of type $\tau \triangleright \sigma$ to retype the term M from type τ to σ . Types are extended with type variables α , lower bounded polymorphic types $\forall(\alpha \triangleleft \tau) \tau$, and the bottom type \perp . A syntactic class for instantiation proofs is added and written with the letter G . They contain bottom instantiations $\perp\tau$, variables c_α , inside instantiations $\forall(\triangleleft G)$, under instantiations $\forall(\alpha \triangleleft c :) G$, elimination instantiations $\&$, introduction instantiations \mathfrak{A} , and the usual transitivity $G \circ G$ and reflexivity \diamond . Finally, we extend environments with type and instance bindings. Type bindings are similar to System F , while instance bindings are new. In the implicit version, we remember that type σ is an instance of type τ by writing $(\tau \triangleright \sigma)$. However, for the explicit version, we need to give a name to the proof that σ is an instance of τ , so we write $(c_\alpha : \tau \triangleright \sigma)$. All these changes are summed up on Figure 3.25.

Evaluation contexts, prevalues, and values for MLF are given in Figure 3.26. Prevalues are extended with instance applications to a prevalue of the elimination instantiation, the under instantiation, the inside instantiation, the bottom instantiation, and instantiation applications to a value of a instance variable. Values are extended with lower bounded abstraction of a value.

Reduction rules for MLF are given on Figure 3.27. The first four rules are those of the STLC. The next rules are new: one for each instantiation proof. Notice however that rules REDREFL and REDTRANS are exactly like those of System F_η . All the new rules are ι -reduction rules since they do not modify the computational content of the term to which they apply.

Rule REDINTRO is similar in spirit to rule REDTLAM of System F_η . However, in System F_η , the type variable is syntactically present, while in MLF it has to be generated. In both cases the variable has to be fresh. In MLF it is generated fresh, while in System F_η it is fresh by typing. So an abstraction with a bottom lower bound can be freely added to a term.

Rule REDELIM is again similar to the rule REDFOR of System F_η : both are instantiation rules. However, there are two main differences. In System F_η , the instantiation argument is

$$\begin{array}{c}
\text{REDCTX} \\
\frac{M \rightsquigarrow_{\beta} N}{E[M] \rightsquigarrow_{\beta} E[N]} \\
\\
\text{REDREFL} \\
\langle \diamond \rangle M \rightsquigarrow_{\iota} M \\
\\
\text{REDAPP} \\
(\lambda(x : \tau) M) N \rightsquigarrow_{\beta} M[x/N] \\
\\
\text{REDFST} \\
\text{fst} \langle M, N \rangle \rightsquigarrow_{\beta} M \\
\\
\text{REDSND} \\
\text{snd} \langle M, N \rangle \rightsquigarrow_{\beta} N \\
\\
\text{REDTRANS} \\
(G_2 \circ G_1) \langle M \rangle \rightsquigarrow_{\iota} G_2 \langle G_1 \langle M \rangle \rangle \\
\\
\text{REDINTRO} \\
\frac{c_{\alpha}, \alpha \notin \text{fv}(M)}{\wp \langle M \rangle \rightsquigarrow_{\iota} \Lambda(\alpha \triangleleft c : \perp) M} \\
\\
\text{REDELIM} \\
&\& \langle \Lambda(\alpha \triangleleft c : \tau) M \rangle \rightsquigarrow_{\iota} M[\alpha/\tau][c_{\alpha}/\diamond] \\
\\
\text{REDUNDER} \\
(\forall(\alpha \triangleleft c : \cdot) G) \langle \Lambda(\alpha \triangleleft c : \tau) M \rangle \rightsquigarrow_{\iota} \Lambda(\alpha \triangleleft c : \tau) G \langle M \rangle \\
\\
\text{REDINSIDE} \\
(\forall(\triangleleft G)) \langle \Lambda(\alpha \triangleleft c : \tau) M \rangle \rightsquigarrow_{\iota} \Lambda(\alpha \triangleleft c : G(\tau)) M[c_{\alpha}/c_{\alpha} \circ G]
\end{array}$$

Figure 3.27: MLF reduction rules

$$\begin{array}{ll}
c_{\alpha}(\tau) = \alpha & \wp(\tau) = \forall(\alpha \triangleleft \perp) \tau \text{ with } \alpha \notin \text{fv}(\tau) \\
(\perp \tau)(\perp) = \tau & \&(\forall(\alpha \triangleleft \tau) \sigma) = \sigma[\alpha/\tau] \\
\diamond(\tau) = \tau & (\forall(\triangleleft G))(\forall(\alpha \triangleleft \tau) \sigma) = \forall(\alpha \triangleleft G(\tau)) \sigma \\
(G_2 \circ G_1)(\tau) = G_2(G_1(\tau)) & (\forall(\alpha \triangleleft c : \cdot) G)(\forall(\alpha \triangleleft \tau) \sigma) = \forall(\alpha \triangleleft \tau) G(\sigma)
\end{array}$$

Figure 3.28: MLF update function

given separately, while in MLF it defaults to the bound τ . The second difference is that the instance variable has also to be substituted in MLF. Since the instance variable is taken equal to the bound, the reflexivity proof witnesses that τ is an instance of τ .

Rule REDUNDER is very close to rule REDCONGR of System F_{η} : both are congruence rules for the polymorphic type. The only difference is anecdotal: MLF handles a lower bound while System F_{η} does not. Notice that, similarly to System F_{η} , the two instance variables have to be the same and the two type variables have to be the same too. If this is not the case, one of them must be renamed before the reduction can occur. The first two variables bind in the instance proof G , while the last two variables bind in the term M .

Finally, rule REDINSIDE has no counter-part in System F_{η} since it is about the bound. If a lower bound abstraction $\Lambda(\alpha \triangleleft c : \tau)M$ encounters an inside instantiation $\forall(\triangleleft G)$, the instantiation proof G modifies the lower bound according to the previous bound τ , written $G(\tau)$ and define in Figure 3.28, and the instance variable c_{α} is updated in order to apply the instantiation G before itself. Together with rule REDELIM it allows arbitrary instantiations with any instance of the lower bound, since it applies a first instantiation followed by an application of the resulting bound.

The term judgment relation is defined on Figure 3.29. The first six rules are those of the STLC. The last two rules are new. However the last one, named TERMTAPP, resembles the rule TERMCONT of System F_{η} . It explains how an instantiation is used: when a term a has type τ with proof M and σ is an instance of τ with proof G , then the term a has also type σ with proof $G \langle M \rangle$. All of this happens under the environment Γ because no bindings occur.

Rule TERMTABS is the abstraction rule for lower bounded polymorphism. If a term a has type σ under the environment Γ extended with the type variable α and its associated

$$\begin{array}{c}
\text{TERMVAR} \\
\frac{\Gamma \text{ env} \quad (x : \tau) \in \Gamma}{x \Rightarrow \Gamma \vdash x : \tau} \\
\\
\text{TERMLAM} \\
\frac{M \Rightarrow \Gamma, (x : \tau) \vdash a : \sigma}{\lambda(x : \tau) M \Rightarrow \Gamma \vdash \lambda x a : \tau \rightarrow \sigma} \\
\\
\text{TERMAPP} \\
\frac{M \Rightarrow \Gamma \vdash a : \tau \rightarrow \sigma \quad N \Rightarrow \Gamma \vdash b : \tau}{MN \Rightarrow \Gamma \vdash ab : \sigma} \\
\\
\text{TERMPAIR} \\
\frac{M \Rightarrow \Gamma \vdash a : \tau \quad N \Rightarrow \Gamma \vdash b : \sigma}{\langle M, N \rangle \Rightarrow \Gamma \vdash \langle a, b \rangle : \tau \times \sigma} \\
\\
\text{TERMFST} \\
\frac{M \Rightarrow \Gamma \vdash a : \tau \times \sigma}{\text{fst } M \Rightarrow \Gamma \vdash \text{fst } a : \tau} \\
\\
\text{TERMSND} \\
\frac{M \Rightarrow \Gamma \vdash a : \tau \times \sigma}{\text{snd } M \Rightarrow \Gamma \vdash \text{snd } a : \sigma} \\
\\
\text{TERMTABS} \\
\frac{M \Rightarrow \Gamma, \alpha, (c_\alpha : \tau \triangleright \alpha) \vdash a : \sigma}{\Lambda(\alpha \triangleleft c : \tau) M \Rightarrow \Gamma \vdash a : \forall(\alpha \triangleleft \tau) \sigma} \\
\\
\text{TERMTAPP} \\
\frac{M \Rightarrow \Gamma \vdash a : \tau \quad G \Rightarrow \Gamma \vdash \tau \triangleright \sigma}{G \langle M \rangle \Rightarrow \Gamma \vdash a : \sigma}
\end{array}$$

Figure 3.29: MLF term judgment relation

$$\begin{array}{c}
\text{INSTBOT} \\
\frac{\Gamma \vdash \tau \text{ type}}{\perp \tau \Rightarrow \Gamma \vdash \perp \triangleright \tau} \\
\\
\text{INSTUNDER} \\
\frac{G \Rightarrow \Gamma, \alpha, (c_\alpha : \tau \triangleright \alpha) \vdash \sigma_1 \triangleright \sigma_2}{\forall(\alpha \triangleleft c : \cdot) G \Rightarrow \Gamma \vdash \forall(\alpha \triangleleft \tau) \sigma_1 \triangleright \forall(\alpha \triangleleft \tau) \sigma_2} \\
\\
\text{INSTABSTR} \\
\frac{(c_\alpha : \tau \triangleright \alpha) \in \Gamma}{c_\alpha \Rightarrow \Gamma \vdash \tau \triangleright \alpha} \\
\\
\text{INSTINSIDE} \\
\frac{G \Rightarrow \Gamma \vdash \tau_1 \triangleright \tau_2}{\forall(\triangleleft G) \Rightarrow \Gamma \vdash \forall(\alpha \triangleleft \tau_1) \sigma \triangleright \forall(\alpha \triangleleft \tau_2) \sigma} \\
\\
\text{INSTINTRO} \\
\frac{\alpha \notin \text{dom}(\Gamma)}{\exists \Rightarrow \Gamma \vdash \tau \triangleright \forall(\alpha \triangleleft \perp) \tau} \\
\\
\text{INSTCOMP} \\
\frac{G_1 \Rightarrow \Gamma \vdash \tau_1 \triangleright \tau_2 \quad G_2 \Rightarrow \Gamma \vdash \tau_2 \triangleright \tau_3}{G_2 \circ G_1 \Rightarrow \Gamma \vdash \tau_1 \triangleright \tau_3} \\
\\
\text{INSTEMELIM} \\
\& \Rightarrow \Gamma \vdash \forall(\alpha \triangleleft \tau) \sigma \triangleright \sigma[\alpha/\tau] \\
\\
\text{INSTID} \\
\Diamond \Rightarrow \Gamma \vdash \tau \triangleright \tau
\end{array}$$

Figure 3.30: MLF instance judgment relation

instance variable c_α witnessing that α is an instance of τ , then it also has type $\forall(\alpha \triangleleft \tau) \sigma$ under environment Γ . The term witnessing this rule is $\Lambda(\alpha \triangleleft c : \tau)M$.

The instance relation of MLF is similar in some points to the containment relation of System F_η . Its judgment is written $G \Rightarrow \Gamma \vdash \tau \triangleright \sigma$, as in System F_η , and means that G is a proof that the type σ is an instance of the type τ under the environment Γ . The rules are given on Figure 3.30. Similarly to System F_η , the extraction property for this judgment takes the well-formedness of its left-hand type as an input. Rules `INSTID` and `INSTCOMP` are exactly the rules `CONTRFL` and `CONSTRANS` of System F_η . Rule `INSTINTRO` is very similar to rule `CONTLAM` of System F_η since it generalizes over a free type variable. The difference is about the bound, but since the bound is the bottom type, it behaves exactly as a polymorphic type.

Rule `INSTEMELIM` resembles rule `CONTTAPP` of System F_η since it is about instantiation. The difference is that the instance type argument is not provided but defaults to the lower bound τ . We see in this rule that MLF does not allow recursive bounds. With recursive bounds α

$$\begin{array}{c}
\text{TYPEVAR} \\
\frac{\Gamma \text{ env} \quad \alpha \in \Gamma}{\Gamma \vdash \alpha \text{ type}} \\
\\
\text{TYPEARR} \\
\frac{\Gamma \vdash \tau \text{ type} \quad \Gamma \vdash \sigma \text{ type}}{\Gamma \vdash \tau \rightarrow \sigma \text{ type}} \\
\\
\text{TYPEPROD} \\
\frac{\Gamma \vdash \tau \text{ type} \quad \Gamma \vdash \sigma \text{ type}}{\Gamma \vdash \tau \times \sigma \text{ type}} \\
\\
\text{TYPEFOR} \\
\frac{\Gamma \vdash \tau \text{ type} \quad \Gamma, \alpha, (c_\alpha : \tau \triangleright \alpha) \vdash \sigma \text{ type}}{\Gamma \vdash \forall(\alpha \triangleleft \tau) \sigma \text{ type}} \\
\\
\text{TYPEBOT} \\
\frac{\Gamma \text{ env}}{\Gamma \vdash \perp \text{ type}} \\
\\
\text{ENVMPTY} \\
\emptyset \text{ env} \\
\\
\text{ENVTERM} \\
\frac{x \notin \text{dom}(\Gamma) \quad \Gamma \vdash \tau \text{ type}}{\Gamma, (x : \tau) \text{ env}} \\
\\
\text{ENVTYPE} \\
\frac{\alpha, c_\alpha \notin \text{dom}(\Gamma) \quad \Gamma \vdash \tau \text{ type}}{\Gamma, \alpha, (c_\alpha : \tau \triangleright \alpha) \text{ env}}
\end{array}$$

Figure 3.31: MLF well-formedness relations

would be free in τ and the substitution $\sigma[\alpha/\tau]$ would be ill-formed. With recursive bounds the solution would be to instantiate with the recursive type $\mu\alpha\tau$ and the proof that $\mu\alpha\tau$ is an instance of $\tau[\alpha/\mu\alpha\tau]$ would be the unfolding coercion (see rule `COERUNFOLD` in Figure 5.6).

Rule `INSTABSTR` looks for an instance hypothesis in the environment. If the type variable α is an instance of the type τ by hypothesis in environment Γ with name c_α , then c_α is a proof that α is an instance of τ under Γ . By rule `INSTBOT`, any well-formed type τ is an instance of the bottom type \perp with proof $\perp\tau$.

Rule `INSTUNDER` is the body congruence rule for the lower bounded polymorphic type, while rule `INSTINSIDE` is the bound congruence rule. A lower bounded polymorphic type is an instance of another lower bounded polymorphic type if they have the same bound and the body of the first is an instance of the body of the second under the same environment extended with the assumption of the mutual bound. Said otherwise: if the type σ_2 is an instance of type σ_1 under the environment Γ extended with the type variable α and the assumption that α is an instance of τ , then the lower bounded polymorphic type $\forall(\alpha \triangleleft \tau) \sigma_2$ is an instance of the lower bounded polymorphic type $\forall(\alpha \triangleleft \tau) \sigma_1$. Similarly for the bound congruence, two lower bounded polymorphic types are in the instance relation if their bounds are also in the instance relation in the same order. If the bound τ_2 is an instance of the bound τ_1 under the environment Γ (because recursive bounds are not allowed), then the lower bounded polymorphic type $\forall(\alpha \triangleleft \tau_2) \sigma$ is an instance of the lower bounded polymorphic type $\forall(\alpha \triangleleft \tau_1) \sigma$ under the same environment Γ .

Finally the well-formedness relations are given on Figure 3.31. The first three rules for type well-formedness are the same as those for System F. The next two rules are new. Rule `TYPEFOR` is an adaptation of the System F rule of the same name. A lower bounded polymorphic type is well-formed if its body type is well-formed under the environment extended with its binding, namely the type variable α and the instance assumption that α is an instance of the lower well-formed bound τ . Rule `TYPEBOT` tells that the bottom type is well-formed under any well-formed environment.

The first two rules of the environment well-formedness are the same as those of the STLC. The last rule, namely `ENVTYPE`, is new. The extended environment $\Gamma, \alpha, (c_\alpha : \tau \triangleright \alpha)$ is well-formed if τ is well-formed under the environment Γ . Besides, the instance variable c_α and the type variable α must not already be bound in Γ .

α, β	Type variables
c	Subtyping variables
$M, N ::= x \mid \lambda(x : \tau) M \mid M M \mid \langle M, M \rangle \mid \text{fst } M \mid \text{snd } M$ $\mid \Lambda(\alpha \triangleright c : \tau) M \mid M[\tau \triangleright G] \mid G\langle M \rangle$	Explicit terms
$\tau, \sigma, \rho ::= \alpha \mid \tau \rightarrow \tau \mid \tau \times \tau \mid \forall(\alpha \triangleright \tau) \tau \mid \top$	Types
$G ::= \diamond \mid G \circ G \mid c \mid \top \mid G \xrightarrow{\tau} G \mid G \times G \mid \forall(\alpha \triangleright c : \tau)[G, G]$	Explicit subtypings
$\Gamma ::= \emptyset \mid \Gamma, (x : \tau) \mid \Gamma, \alpha \mid \Gamma, (c : \tau \triangleright \tau)$	Environments

Figure 3.32: System $F_{<}$: syntax

3.5.2 Properties

The properties of MLF are similar to those of System F with some modifications about the bounds. The type substitution lemma of System F is replaced by a lower bound substitution lemma.

Lemma 41 (Lower bound substitution). *If $\Gamma \vdash \sigma$ type and $G \Rightarrow \Gamma \vdash \tau \triangleright \sigma$ hold, then the following assertions hold.*

- If $\Gamma, \alpha, (c_\alpha : \tau \triangleright \alpha), \Gamma' \vdash \rho$ type holds, then $\Gamma, \Gamma'[\alpha/\sigma] \vdash \rho[\alpha/\sigma]$ type holds.
- If $M \Rightarrow \Gamma, \alpha, (c_\alpha : \tau \triangleright \alpha), \Gamma' \vdash a : \rho$ holds, then $M[\alpha/\sigma][c_\alpha/G] \Rightarrow \Gamma, \Gamma'[\alpha/\sigma] \vdash a : \rho[\alpha/\sigma]$ holds.
- If $G' \Rightarrow \Gamma, \alpha, (c_\alpha : \tau \triangleright \alpha), \Gamma' \vdash \rho \triangleright \rho'$ holds, then $G'[\alpha/\sigma][c_\alpha/G] \Rightarrow \Gamma, \Gamma'[\alpha/\sigma] \vdash \rho[\alpha/\sigma] \triangleright \rho'[\alpha/\sigma]$ holds.

3.6 System $F_{<}$:

System $F_{<}$ is another extension of System F dealing with the notion of subtyping. However, there are two main differences with System F_η . On the one hand, System $F_{<}$ lacks the distributivity rule. However on the other hand, System $F_{<}$ enjoys upper bounded polymorphism, which can be seen as an analog of the lower bounded polymorphism of MLF. The analogy only holds for the syntax of types: lower bounded polymorphism in MLF is fully composable, while upper bounded polymorphism is shallow in $F_{<}$. We detail these ideas with coercions in Section 4.5.4. The polymorphic type of System $F_{<}$ abstracts over a type variable smaller than an upper bound type. We recover the polymorphic type of System F by using a particular upper bound called top. All types are smaller than the top type, and thus we can instantiate an upper bounded polymorphic type with bound top with any type we want as we can already do in System F. This is why System $F_{<}$ is an extension of System F in terms of features. But we actually prefer to see it as a syntactical extension of the STLC with subtyping, the top type, and upper bounded polymorphism. We will see in Section 4.5.4 that System $F_{<}$ is actually not complete according to how it extends the STLC.

3.6.1 Definition

System $F_{<}$ extends the STLC syntax with type variables, written α or β , and subtyping variables, written c . Subtyping variables are only necessary in the explicit version of the

$ \begin{aligned} E ::= & \lambda(x : \tau) \square \mid \square \mid M \square \mid \langle \square, M \rangle \mid \langle M, \square \rangle \mid \mathbf{fst} \square \mid \mathbf{snd} \square \\ & \mid \Lambda(\alpha \triangleright c : \tau) \square \mid \square[\tau \triangleright G] \mid G(\square) \\ p ::= & x \mid p v \mid \mathbf{fst} p \mid \mathbf{snd} p \mid p[\tau \triangleright G] \mid (\forall(\alpha \triangleright c : \tau)[G, G])\langle p \rangle \\ & \mid (G \xrightarrow{\tau} G)\langle p \rangle \mid (G \times G)\langle p \rangle \mid c\langle v \rangle \\ v ::= & p \mid \lambda(x : \tau) v \mid \langle v, v \rangle \mid \Lambda(\alpha \triangleright c : \tau) v \mid \top\langle v \rangle \end{aligned} $	<p>Evaluation contexts</p> <p>Prevalues</p> <p>Values</p>
--	---

Figure 3.33: System $F_{<}$: notations

$ \begin{array}{c} \text{REDCTX} \\ \frac{M \rightsquigarrow_{\beta} N}{E[M] \rightsquigarrow_{\beta} E[N]} \end{array} $	$ \begin{array}{c} \text{REDAPP} \\ (\lambda(x : \tau) M) N \rightsquigarrow_{\beta} M[x/N] \end{array} $	$ \begin{array}{c} \text{REDFST} \\ \mathbf{fst} \langle M, N \rangle \rightsquigarrow_{\beta} M \end{array} $	$ \begin{array}{c} \text{REDSND} \\ \mathbf{snd} \langle M, N \rangle \rightsquigarrow_{\beta} N \end{array} $
$ \begin{array}{c} \text{REDFOR} \\ (\Lambda(\alpha \triangleright c : \tau) M)[\sigma \triangleright G] \rightsquigarrow_{\iota} M[\alpha/\sigma][c/G] \end{array} $	$ \begin{array}{c} \text{REDREFL} \\ \diamond \langle M \rangle \rightsquigarrow_{\iota} M \end{array} $	$ \begin{array}{c} \text{REDTRANS} \\ (G_2 \circ G_1)\langle M \rangle \rightsquigarrow_{\iota} G_2\langle G_1\langle M \rangle \rangle \end{array} $	
$ \begin{array}{c} \text{REDARR} \\ (G_1 \xrightarrow{\tau} G_2)\langle \lambda(x : \tau') M \rangle \rightsquigarrow_{\iota} \lambda(x : \tau) G_2\langle M[x/G_1\langle x \rangle] \rangle \end{array} $			
$ \begin{array}{c} \text{REDPROD} \\ (G_1 \times G_2)\langle \langle M_1, M_2 \rangle \rangle \rightsquigarrow_{\iota} \langle G_1\langle M_1 \rangle, G_2\langle M_2 \rangle \rangle \end{array} $			
$ \begin{array}{c} \text{REDCONGR} \\ (\forall(\alpha \triangleright c : \tau')[G_1, G_2])\langle \Lambda(\alpha \triangleright c : \tau) M \rangle \rightsquigarrow_{\iota} \Lambda(\alpha \triangleright c : \tau') G_2\langle M[c/G_1] \rangle \end{array} $			

Figure 3.34: System $F_{<}$: reduction rules

type system. Types are extended with type variables α , upper bounded polymorphic types $\forall(\alpha \triangleright \tau) \tau$, and the top type \top . Terms are extended with upper bounded type abstractions $\Lambda(\alpha \triangleright c : \tau) M$, upper bounded type applications $M[\tau \triangleright G]$, and subtyping applications $G\langle M \rangle$. A new class for subtypings is added. Subtypings are written G and contain reflexivity \diamond , transitivity $G \circ G$, variables c , top subtypings \top , arrow congruence $G \xrightarrow{\tau} G$, product congruence $G \times G$, and upper bounded congruence $\forall(\alpha \triangleright c : \tau)[G, G]$. Finally, we extend environments with type and subtyping bindings. Type bindings are similar to System F , while subtyping bindings are new. In the implicit version, we remember that type τ is a subtype of type σ by writing $(\tau \triangleright \sigma)$. However, for the explicit version, we need to give a name to the proof that τ is a subtype of σ , so we write $(c : \tau \triangleright \sigma)$. These modifications can be found on Figure 3.32.

Notations for System $F_{<}$ follow the same logic as the STLC. They are given on Figure 3.33. Evaluation contexts are all one-hole terms of depth one and correspond to strong reduction. Prevalues are variables or destructors applied to values but where a constructor is expected, in which case only a prevalue is applied. And values are prevalues or constructors applied to values. The new prevalues are: upper bounded application since it waits an upper bounded abstraction, congruences because they wait their constructor (abstraction or pair), and subtyping variables applied to a value. The new values are upper bounded abstraction, and the top subtyping.

Reduction rules for System $F_{<}$ are given on Figure 3.34. The first four rules are those of

$$\begin{array}{c}
\text{TERMVAR} \\
\frac{\Gamma \text{ env} \quad (x : \tau) \in \Gamma}{x \Rightarrow \Gamma \vdash x : \tau} \\
\\
\text{TERMLAM} \\
\frac{M \Rightarrow \Gamma, (x : \tau) \vdash a : \sigma}{\lambda(x : \tau) M \Rightarrow \Gamma \vdash \lambda x a : \tau \rightarrow \sigma} \\
\\
\text{TERMAPP} \\
\frac{M \Rightarrow \Gamma \vdash a : \tau \rightarrow \sigma \quad N \Rightarrow \Gamma \vdash b : \tau}{MN \Rightarrow \Gamma \vdash ab : \sigma} \\
\\
\text{TERMPAIR} \\
\frac{M \Rightarrow \Gamma \vdash a : \tau \quad N \Rightarrow \Gamma \vdash b : \sigma}{\langle M, N \rangle \Rightarrow \Gamma \vdash \langle a, b \rangle : \tau \times \sigma} \\
\\
\text{TERMFST} \\
\frac{M \Rightarrow \Gamma \vdash a : \tau \times \sigma}{\text{fst } M \Rightarrow \Gamma \vdash \text{fst } a : \tau} \\
\\
\text{TERMSND} \\
\frac{M \Rightarrow \Gamma \vdash a : \tau \times \sigma}{\text{snd } M \Rightarrow \Gamma \vdash \text{snd } a : \sigma} \\
\\
\text{TERMGEN} \\
\frac{M \Rightarrow \Gamma, \alpha, (c : \alpha \triangleright \tau) \vdash a : \rho}{\Lambda(\alpha \triangleright c : \tau) M \Rightarrow \Gamma \vdash a : \forall(\alpha \triangleright \tau) \rho} \\
\\
\text{TERMINST} \\
\frac{M \Rightarrow \Gamma \vdash a : \forall(\alpha \triangleright \tau) \rho \quad \Gamma \vdash \sigma \text{ type} \quad G \Rightarrow \Gamma \vdash \sigma \triangleright \tau[\alpha/\sigma]}{M[\sigma \triangleright G] \Rightarrow \Gamma \vdash a : \rho[\alpha/\sigma]} \\
\\
\text{TERMSUB} \\
\frac{M \Rightarrow \Gamma \vdash a : \tau \quad G \Rightarrow \Gamma \vdash \tau \triangleright \sigma}{G\langle M \rangle \Rightarrow \Gamma \vdash a : \sigma}
\end{array}$$

Figure 3.35: System $F_{<}$: term judgment relation

the STLC. The next rules are new: one for each subtyping proof. Notice however that rules REDREFL, REDTRANS, REDARR, and REDPROD are exactly like those of System F_{η} . So we only focus on the two new rules REDFOR and REDCONGR, which are actually ameliorations of the System F_{η} rules of the same name. All these rules are ι -reduction rules since they do not modify the computational content of the term on which they apply.

Rule REDFOR matches an upper bounded application over an upper bounded abstraction. This is a regular scheme for abstraction constructs. It reduces to a substitution. In the current case, we not only substitutes the type variable α with the argument type σ as it is done in the System F rule, but we also substitute the subtyping variable c with the argument subtyping proof G .

Finally, rule REDCONGR applies a subtyping proof under an upper bounded abstraction. This is more technical than in System F_{η} since now the bounds may be modified too. The binding $(c : \alpha \triangleright \tau')$ is the new binding, G_2 is the subproof to apply on the term body and G_1 is the subtyping proof to update the bound from τ to τ' and replaces the old subtyping variable c .

The term judgment relation is defined on Figure 3.35. The first six rules are those of the STLC. The last three rules are new. However the last one, named TERMSUB, resembles the rule TERMCONT of System F_{η} . It explains how a subtyping proof is used: when a term a has type τ with proof M and τ is a subtype of σ with proof G , then the term a has also type σ with proof $G\langle M \rangle$. All of this happens under the environment Γ because no bindings occur.

Rules TERMGEN and TERMINST are ameliorations of the System F rules of the same name. The abstraction and application are now for upper bounded types. If a term a has type ρ under an extended environment $\Gamma, \alpha, (c : \alpha \triangleright \tau)$, where c is a subtyping proof that α is a subtype of τ , with witness M , then it also has type $\forall(\alpha \triangleright \tau) \rho$ under environment Γ with

$$\begin{array}{c}
\text{SUBREFL} \\
\diamond \Rightarrow \Gamma \vdash \tau \triangleright \tau \\
\\
\text{SUBTRANS} \\
\frac{G_1 \Rightarrow \Gamma \vdash \tau_1 \triangleright \tau_2 \quad G_2 \Rightarrow \Gamma \vdash \tau_2 \triangleright \tau_3}{G_2 \circ G_1 \Rightarrow \Gamma \vdash \tau_1 \triangleright \tau_3} \\
\\
\text{SUBVAR} \\
\frac{(c : \tau \triangleright \sigma) \in \Gamma}{c \Rightarrow \Gamma \vdash \tau \triangleright \sigma} \\
\\
\text{SUBTOP} \\
\top \Rightarrow \Gamma \vdash \tau \triangleright \top \\
\\
\text{SUBARR} \\
\frac{\Gamma \vdash \tau \text{ type} \quad G_1 \Rightarrow \Gamma \vdash \tau \triangleright \tau' \quad G_2 \Rightarrow \Gamma \vdash \sigma' \triangleright \sigma}{G_1 \xrightarrow{\tau} G_2 \Rightarrow \Gamma \vdash \tau' \rightarrow \sigma' \triangleright \tau \rightarrow \sigma} \\
\\
\text{SUBPROD} \\
\frac{G_1 \Rightarrow \Gamma \vdash \tau' \triangleright \tau \quad G_2 \Rightarrow \Gamma \vdash \sigma' \triangleright \sigma}{G_1 \times G_2 \Rightarrow \Gamma \vdash \tau' \times \sigma' \triangleright \tau \times \sigma} \\
\\
\text{SUBCONGR} \\
\frac{\Gamma, \alpha \vdash \tau' \text{ type} \quad G_1 \Rightarrow \Gamma, \alpha, (c : \alpha \triangleright \tau') \vdash \alpha \triangleright \tau \quad G_2 \Rightarrow \Gamma, \alpha, (c : \alpha \triangleright \tau') \vdash \sigma \triangleright \sigma'}{\forall(\alpha \triangleright c : \tau')[G_1, G_2] \Rightarrow \Gamma \vdash \forall(\alpha \triangleright \tau) \sigma \triangleright \forall(\alpha \triangleright \tau') \sigma'}
\end{array}$$

Figure 3.36: System $F_{<}$: subtyping judgment relation

witness $\Lambda(\alpha \triangleright c : \tau) M$. This witness binds c in M and α in τ and M . As a consequence, unlike for MLF, upper bounds are recursive in System $F_{<}$. If a term a has type $\forall(\alpha \triangleright \tau) \rho$ under environment Γ with witness M , then it also has type $\rho[\alpha/\sigma]$ under environment Γ by instantiation with the well-formed type σ and the proof G that σ is a subtype of $\tau[\alpha/\sigma]$. The type variable α may appear in the type τ , since we allow recursive type. This is why the right-hand type of the subtyping proof is not simply τ .

The subtyping relation of System $F_{<}$ is similar in some points to the containment relation of System F_η . Its judgment is written $G \Rightarrow \Gamma \vdash \tau \triangleright \sigma$, as in System F_η , and means that G is a proof that the type τ is a subtype of the type σ under the environment Γ . The rules are given on Figure 3.36. Similarly to System F_η , the extraction property for this judgment takes the well-formedness of its left-hand type as an input. Rule SUBREFL, SUBTRANS, SUBARR, and SUBPROD are exactly the rules CONTREFL, CONTTRANS, CONTARR, and CONTPROD of System F_η .

Rule SUBVAR looks a subtyping hypothesis in the environment. If the type τ is a subtype of the type σ by hypothesis in environment Γ with name c , then c is a proof that τ is a subtype of σ under Γ . By rule SUBTOP, any type τ is a subtype of the top type \top with proof \top .

Rule SUBCONGR is a congruence rule for the upper bounded polymorphic type. It is similar to rule CONTCONGR of System F_η , but it is more technical and complicated since it has to handle the bound. If, assuming that type variable α is a subtype of type τ' , we can show that it is also a subtype of τ and we can also show that type σ is a subtype of σ' then we have that $\forall(\alpha \triangleright \tau) \sigma$ is a subtype of $\forall(\alpha \triangleright \tau') \sigma'$. The first premise about the well-formedness of τ' is necessary to prove the extraction lemma. We can wonder why we need this particular subtyping premise for the bound. The body subtyping seems natural. Rule SUBCONGR actually differs depending on the variant of System $F_{<}$ we consider. We present here two other variants:

$$\begin{array}{c}
\text{KERNEL-FSUB} \\
\frac{\Gamma, \alpha, (\alpha \triangleright \tau) \vdash \sigma \triangleright \sigma'}{\Gamma \vdash \forall(\alpha \triangleright \tau) \sigma \triangleright \forall(\alpha \triangleright \tau) \sigma'} \\
\\
\text{FULL-FSUB} \\
\frac{\Gamma \vdash \tau' \triangleright \tau \quad \Gamma, \alpha, (\alpha \triangleright \tau') \vdash \sigma \triangleright \sigma'}{\Gamma \vdash \forall(\alpha \triangleright \tau) \sigma \triangleright \forall(\alpha \triangleright \tau') \sigma'}
\end{array}$$

$$\begin{array}{c}
\text{TYPEVAR} \\
\frac{\Gamma \text{ env} \quad \alpha \in \Gamma}{\Gamma \vdash \alpha \text{ type}} \\
\\
\text{TYPEARR} \\
\frac{\Gamma \vdash \tau \text{ type} \quad \Gamma \vdash \sigma \text{ type}}{\Gamma \vdash \tau \rightarrow \sigma \text{ type}} \\
\\
\text{TYPEPROD} \\
\frac{\Gamma \vdash \tau \text{ type} \quad \Gamma \vdash \sigma \text{ type}}{\Gamma \vdash \tau \times \sigma \text{ type}} \\
\\
\text{TYPEFOR} \\
\frac{\Gamma, \alpha \vdash \tau \text{ type} \quad \Gamma, \alpha, (c : \alpha \triangleright \tau) \vdash \sigma \text{ type}}{\Gamma \vdash \forall(\alpha \triangleright \tau) \sigma \text{ type}} \\
\\
\text{TOTYPE} \\
\frac{\Gamma \text{ env}}{\Gamma \vdash \top \text{ type}} \\
\\
\text{ENVEMPTY} \\
\frac{}{\emptyset \text{ env}} \\
\\
\text{ENVTERM} \\
\frac{x \notin \text{dom}(\Gamma) \quad \Gamma \vdash \tau \text{ type}}{\Gamma, (x : \tau) \text{ env}} \\
\\
\text{ENVTYPE} \\
\frac{\alpha, c \notin \text{dom}(\Gamma) \quad \Gamma, \alpha \vdash \tau \text{ type}}{\Gamma, \alpha, (c : \alpha \triangleright \tau) \text{ env}}
\end{array}$$

Figure 3.37: System $F_{<}$: well-formedness relations

In Kernel $F_{<}$, the bounds τ and τ' must be equal, whereas Full $F_{<}$ only requires the bound τ' to be a subtype of the bound τ . Moreover, α cannot appear free in the bounds τ or τ' in Kernel or Full $F_{<}$, while the presented variant, called F -bounded, allows this form of recursion. The most general assumption, $\Gamma, \alpha, (\alpha \triangleright \tau') \vdash \alpha \triangleright \tau$ is that of F -bounded. Perhaps surprisingly, this is a slightly more general rule [5] than the more intuitive one $\Gamma, \alpha, (\alpha \triangleright \tau') \vdash \tau' \triangleright \tau$. In summary, we have $\text{Kernel } F_{<} \subset \text{Full } F_{<} \subset F\text{-bounded}$ where all inclusions are strict.

Finally the well-formedness relations are given on Figure 3.37. The first three rules for type well-formedness are the same as those for System F . The next two rules are new. Rule TYPEFOR is an adaptation of the System F rule of the same name. An upper bounded polymorphic type is well-formed if its body type is well-formed under the environment extended with its binding. Rule TOTYPE tells that the top type is well-formed under any well-formed environment. The first two rules of the environment well-formedness are the same as those of the STLC . The last rule, namely ENVTYPE , is new. The extended environment $\Gamma, \alpha, (c : \alpha \triangleright \tau)$ is well-formed if τ is well-formed under environment Γ extended with α . Besides, the subtyping variable c and the type variable α must not be already bound in Γ .

3.6.2 Properties

The properties of System $F_{<}$ are similar to those of System F with some modifications about the bounds. The type substitution lemma of System F is replaced by an upper bound substitution lemma.

Lemma 42 (Upper bound substitution). *If $\Gamma \vdash \sigma \text{ type}$ and $G \Rightarrow \Gamma \vdash \sigma \triangleright \tau[\alpha/\sigma]$ hold, then the following assertions hold.*

- If $\Gamma, \alpha, (c : \alpha \triangleright \tau), \Gamma' \vdash \rho \text{ type}$ holds, then $\Gamma, \Gamma'[\alpha/\sigma] \vdash \rho[\alpha/\sigma] \text{ type}$ holds.
- If $M \Rightarrow \Gamma, \alpha, (c : \alpha \triangleright \tau), \Gamma' \vdash a : \rho$ holds, then $M[\alpha/\sigma][c/G] \Rightarrow \Gamma, \Gamma'[\alpha/\sigma] \vdash a : \rho[\alpha/\sigma]$ holds.
- If $G' \Rightarrow \Gamma, \alpha, (c : \alpha \triangleright \tau), \Gamma' \vdash \rho \triangleright \rho'$ holds, then $G'[\alpha/\sigma][c/G] \Rightarrow \Gamma, \Gamma'[\alpha/\sigma] \vdash \rho[\alpha/\sigma] \triangleright \rho'[\alpha/\sigma]$ holds.

α, β		Type variables
	$\tau ::= \alpha \mid \tau \rightarrow \tau \mid \tau \times \tau$	Types
	$\sigma ::= \tau \mid \forall \bar{\alpha}. C \Rightarrow \sigma$	Type schemes
	$\Gamma ::= \emptyset \mid \Gamma, (x : \sigma)$	Environments

Figure 3.38: Constraint ML syntax

$\frac{\text{TERMVAR}}{(x : \sigma) \in \Gamma}{C; \Gamma \vdash x : \sigma}$	$\frac{\text{TERMSUB}}{C; \Gamma \vdash a : \tau \quad C \vdash \tau \triangleright \tau'}{C; \Gamma \vdash a : \tau'}$	$\frac{\text{TERMABS}}{C; \Gamma, (x : \tau) \vdash a : \tau'}{C; \Gamma \vdash \lambda x a : \tau \rightarrow \tau'}$
$\frac{\text{TERMAPP}}{C; \Gamma \vdash a_1 : \tau_1 \rightarrow \tau_2 \quad C; \Gamma \vdash a_2 : \tau_1}{C; \Gamma \vdash a_1 a_2 : \tau_2}$	$\frac{\text{TERMPAIR}}{C; \Gamma \vdash a_1 : \tau_1 \quad C; \Gamma \vdash a_2 : \tau_2}{C; \Gamma \vdash \langle a_1, a_2 \rangle : \tau_1 \times \tau_2}$	
$\frac{\text{TERMFST}}{C; \Gamma \vdash a : \tau_1 \times \tau_2}{C; \Gamma \vdash \text{fst } a : \tau_1}$	$\frac{\text{TERMSND}}{C; \Gamma \vdash a : \tau_1 \times \tau_2}{C; \Gamma \vdash \text{snd } a : \tau_2}$	$\frac{\text{TERMLET}}{C; \Gamma \vdash a : \sigma \quad C; \Gamma, (x : \sigma) \vdash a' : \tau'}{C; \Gamma \vdash (\lambda x a') a : \tau'}$
$\frac{\text{TERMINTRO}}{C \wedge D; \Gamma \vdash a : \tau \quad \bar{\alpha} \notin \text{fv}(C, \Gamma)}{C \wedge \exists \bar{\alpha}. D; \Gamma \vdash a : \forall \bar{\alpha}. D \Rightarrow \tau}$	$\frac{\text{TERME LIM}}{C; \Gamma \vdash a : \forall \bar{\alpha}. D \Rightarrow \tau' \quad C \vdash D[\bar{\alpha} \leftarrow \tau]}{C; \Gamma \vdash a : \tau'[\bar{\alpha} \leftarrow \tau]}$	

Figure 3.39: Constraint ML term judgment relation

3.7 Constraint ML

Constraint ML is an inference type system with constraints in the tradition of the Hindley/Milner type system, where syntactical types are split between simple types and type schemes. In **Constraint ML**, type schemes are extended to constrained type schemes. The presentation we give is inspired from [27]. The reason to choose this presentation rather than the more general setting of Chapter 10 of [30] is to keep the presentation lighter. We only define the implicit version of the type system since it does not have an explicit version.

The syntax of **Constraint ML** is defined on Figure 3.38. It extends the STLC with type variables α or β . Types contain type variables α , arrow types $\tau \rightarrow \tau$, and product types $\tau \times \tau$. Type schemes contain types τ and constrained type schemes $\forall \bar{\alpha}. C \Rightarrow \sigma$. The constrained type scheme $\forall \bar{\alpha}. C \Rightarrow \sigma$ quantifies over the type variables $\bar{\alpha}$ that satisfy the constraint C , which we handle in an abstract way. The type variables $\bar{\alpha}$ are bound in the constraint C and the body of the type scheme σ . In terms of expressivity, this is the most interesting part of **Constraint ML**: constraint abstraction. Of course the inference part is interesting, but since our main concern are kernel type systems, we only focus on expressivity. Finally, environments are lists of term bindings of the form $(x : \sigma)$ binding a term variable x to its type scheme σ .

The term judgment is written $C; \Gamma \vdash a : \sigma$ meaning that the term a has type scheme σ under environment Γ given the constraint C holds. The rules are given on Figure 3.39. Rule **TERMVAR** looks up in the environment to find the type scheme of a term variable. Rule **TERMSUB** allows to change the type of a term from τ to τ' under environment Γ with respect

to the constraint C , provided $\tau \triangleright \tau'$ is a valid constraint with respect to C . We deliberately leave the notion of constraint abstract, since we want to focus on the mechanisms.

Rule `TERMABS` and `TERMAPP` are the usual typing rules for the arrow type. If a term a has type τ' under an environment extended with $(x : \tau)$ then the term $\lambda x a$ has type $\tau \rightarrow \tau'$. And if the term a_1 has type $\tau_1 \rightarrow \tau_2$ and the term a_2 has type τ_1 , then the application $a_1 a_2$ has type τ_2 . Rules `TERMPAIR`, `TERMFST`, and `TERMSND` are similar to their STLC alternatives.

Rule `TERMLET` explains how to type the let-binding: $(\lambda x a') a$. This rule is not simply derivable from rules `TERMABS` and `TERMAPP` since the term variable x gets associated to a type scheme and not simply a type. This comes from the fact that `Constraint ML` does inference. Since the abstraction is directly applied, its argument type can be generalized.

Rules `TERMINTRO` and `TERMELIM` are the two interesting rules. To generalize a type τ over the type variables $\bar{\alpha}$, the usual condition that the type variables $\bar{\alpha}$ should not be free in the environment Γ has to be satisfied. But this is not enough. The generalized type has to gather enough constraints, denoted as D , in order for the remaining constraint C not to mention the type variables $\bar{\alpha}$. Finally, the conclusion constraint to be satisfied is not simply C , but $C \wedge \exists \bar{\alpha}. D$. This additional condition is present to make sure that the erasable abstraction is inhabited: there is an instantiation for the type variables $\bar{\alpha}$ such that the constraint D holds for this instantiation. This condition is necessary for soundness. Without it, it would be possible to abstract over $\tau_1 \rightarrow \tau_2 \triangleright \sigma_1 \times \sigma_2$ for example. Finally, a quantified type $\forall \bar{\alpha}. D \Rightarrow \tau'$ can be instantiated with types $\bar{\tau}$, as long as the constraint $D[\bar{\alpha} \leftarrow \bar{\tau}]$ holds under the constraint C .

3.8 Existing Coercions

We define coercions as erasable and composable typing transformations. In particular, coercions do not modify the computational content of the term they are applied to. As a consequence, coercions are only visible in the explicit version of type systems. This definition is closely related to subtyping, since subtyping defines a relation on types which is composable and do not affect computational content. Actually, subtyping is a particular case of coercions: those that do not alter the typing environment.

Since coercions are composable typing transformations, we may wonder if it would be possible to have all typing transformations of a type system to be expressed as coercions and thus be composable. This is our approach in Part II. The introduction and elimination rules of erasable types are all expressed as coercions. This naturally gives fully access to the concerned type feature: erasable and composable introduction and elimination rule. The typical counter-example is System $F_{<}$, which features upper bounded polymorphism only partially: it is an erasable type feature, but there is no composable introduction and elimination rule (see Section 4.5.4).

Some existing type systems already use coercions as a tool, even if their whole type system is not expressed as coercions. System F_η , which we described in Section 3.4 uses containments as coercions. System $F_{<}$, which we defined in Section 3.6 uses subtyping as coercions. `MLF`, defined in Section 3.5, uses instantiations as coercions. However, all three type systems are not fully expressed as coercions. In particular, erasable abstractions (introduction rules of erasable types) are not included in the coercion judgment but only accessible to the term judgment. As a consequence, η -expansion is not sufficient for deep generalization, and distributivity rules are necessary, like in System F_η . Erasable abstractions are: type abstraction for System F_η ,

upper bounded abstraction for System $F_{<}$, and lower bounded abstraction for MLF. The case for Constraint ML is worse, because erasable quantifiers are not even part of types, but only accessible in type schemes. So there is no hope to use η -expansion to have deep generalization, because such constraint would have no syntax.

The type systems defined in this chapter have distinct sets of features, but it looks like they rely on more general ideas. As such, we may wonder if we can unify these type systems in a single framework. This factorization would ease syntactical comparison of these type systems and give results about the compatibility and orthogonality of their features (Section 4.5.5). It would also help to understand which features of one type system are missing to another one. For instance, MLF and System $F_{<}$ looks similar and different at the same time and we may wonder whether one can be extended to contain the other one. Both type systems are similar because they feature bounded polymorphism: upper bounded polymorphism for System $F_{<}$ and lower bounded polymorphism for MLF. However, MLF features deep instantiation but misses subtyping, while System $F_{<}$ features subtyping but misses deep instantiation. And finally, the soundness and strong normalization of this unified type system would imply the soundness and strong normalization of all the subsumed type systems. For instance, the strong normalization of MLF has been proved with this approach.

Part II

Type Systems as Coercions

Chapter 4

An explicit calculus of coercions: System F_l^p

System $F_{<}$ and MLF are extensions of System F in somewhat dual ways. They both include bounded polymorphism, which is a form of coercion abstraction: $F_{<}$ has upper bounded polymorphism while MLF has lower bounded polymorphism. However, they have incompatible definitions. The same argument holds for all pairs of existing type systems described in Chapter 3. The goal of this chapter is to define a framework where the features of the type systems described in Chapter 3 are expressed in a uniform way, are orthogonal, and can be freely combined. We shall see in Section 4.6 why this may not be completely possible in the explicit version.

A type system can be seen as the following: a syntax for invariants Φ (what we used to call approximations), typing rules for the language constructs $a : \Phi$, and coercions for invariants $\Phi_1 \triangleright \Phi_2$, which tells us that $a : \Phi_1$ implies $a : \Phi_2$ for all a , or that Φ_1 is a better approximation than Φ_2 . Semantically, the invariants Φ are interpreted as sets of terms, typing rules as proofs of memberships, and coercions as proofs of inclusion. Actually, coercions could be extended to be any kind of proofs and inclusions would just be a particular kind of proofs, see Chapter 5.

Similarly to Chapter 3, we simultaneously define the implicit and explicit version of the type system. According to this view of type systems, the explicit version gives witnesses to the two judgments: the term judgment $a : \Phi$ becomes $M \Rightarrow a : \Phi$ and the coercion judgment $\Phi_1 \triangleright \Phi_2$ becomes $G \Rightarrow \Phi_1 \triangleright \Phi_2$. The term M and the coercion proof G witness the implicit judgments.

As before, we consider strong reduction for two reasons. First, because the soundness of strong reduction implies the soundness of all other reduction strategies. Only the correspondence between syntactic values and semantic values, which is simple to prove, has to be rechecked for each reduction strategy. Then, because strong reduction gives a better insight to the language. Understanding strong reduction is a guarantee that no corner cases have been forgotten. Finally, strong reduction justifies some optimization techniques.

The type system containing all the features of this chapter, called System F_l^p , is given in Section 4.3. The first section describes the bare framework without any additional features and thus corresponds to the STLC. The next section describes extensions as pairwise orthogonal features that can be combined at will. We then describe System F_l^p , give its properties, and discuss its expressiveness. In this framework, invariants Φ are typings of the form $\Gamma \vdash \tau$ where Γ is a list of bindings describing the invariants of the environment in which the term can be

α, β	Type variables
c	Coercion variables
$M, N ::= x \mid \lambda(x : \tau) M \mid M M \mid \langle M, M \rangle \mid \mathbf{fst} M \mid \mathbf{snd} M \mid G \langle M \rangle$	Explicit terms
$\tau, \sigma, \rho ::= \alpha \mid \tau \rightarrow \tau \mid \tau \times \tau$	Types
$G ::= c \mid \diamond \mid G \circ G \mid *G$	Coercions
$\Gamma ::= \emptyset \mid \Gamma, (x : \tau) \mid \Gamma, \alpha \mid \Gamma, (c : \tau \triangleright \tau)$	Environments

Figure 4.1: Base system syntax

$\Sigma ::= \emptyset \mid \Sigma, \alpha \mid \Sigma, (c : \tau \triangleright \tau)$	Erasable environments
$E ::= \lambda(x : \tau) [] \mid [] M \mid M [] \mid \langle [], M \rangle \mid \langle M, [] \rangle \mid \mathbf{fst} [] \mid \mathbf{snd} [] \mid G \langle [] \rangle$	Evaluation contexts
$p ::= x \mid p v \mid \mathbf{fst} p \mid \mathbf{snd} p \mid c \langle v \rangle$	Prevalues
$v ::= p \mid \lambda(x : \tau) v \mid \langle v, v \rangle$	Values

Figure 4.2: Base system notations

considered and τ is a type describing the invariant of the term.

4.1 Base system

The base system for a non-dependent type system¹ on the λ -calculus is the STLC. So, the framework we define is equivalent to the STLC, although its presentation differs a little to better fit in our framework.

The syntax for the base type system is given on Figure 4.1 (we remind the reader that the language definition is given on Figure 2.1). It corresponds to the STLC with two small differences. First, we added variables to all syntactic classes in order to be able to abstract over them later. Notice that we also added the corresponding binders. Type variables are written α or β , while coercion variables are written c . Type bindings are written α and coercion bindings are written $(c : \tau \triangleright \tau)$. Notice that coercion variables only exist in the explicit version: they are always implicit in the implicit version. The second difference is the new syntactical class of *coercions*. Currently, it only contains coercion variables c , weakening $*G$, and operations to close coercions by reflexivity \diamond and transitivity $G \circ G$.

Terms, written M or N , contain variables x , abstractions $\lambda(x : \tau) M$ where the term variable x is annotated with its type τ , applications $M M$, pairs $\langle M, M \rangle$, projections $\text{fst } M$ and $\text{snd } M$, and the new coercion construct $G \langle M \rangle$. This construct is used to change the invariant (called a typing in our framework) of the term M according to the coercion G , interpreted as a proof of inclusion between invariants.

Types, written τ , σ , or ρ , contain variables α , arrow types $\tau \rightarrow \tau$, and product types $\tau \times \tau$. Coercions, written G , contain variables c , reflexivity \diamond , transitivity $G \circ G$, and weakening $*G$. Environments, written Γ , are lists of bindings. Bindings contain term bindings $(x : \tau)$, type bindings α , and coercion bindings $(c : \tau \triangleright \tau)$. The coercion variable in the coercion binding is only necessary in the explicit version of the type system. In the implicit version, coercion bindings can be written $(\tau \triangleright \tau)$, omitting the coercion variable.

We define erasable environments, evaluation contexts, and values in Figure 4.2. We write Σ the subset of environments that are erasable: they contain type and coercion bindings (all bindings but the computational ones: term bindings). Erasable environments are used for typing coercions because coercions may change environments as part of typings.

Evaluation contexts E contain all possible one hole contexts of depth one. We overload the notation from evaluation contexts of the λ -calculus and make the distinction clear when necessary.

Prevalues and values are defined as usual for strong reduction. Prevalues are either variables or destructors applied to prevalues. For instance the application is a destructor expecting a constructor on its left-hand side, so its corresponding prevalue is $p v$. The case for the coercion construct is special since it can be either a constructor or a destructor depending on its left-hand side. And depending on this left-hand side, it will expect or not a constructor on its right-hand side. Reflexivity, transitivity, and coercion variables do not expect a constructor for their coerced term. However, since reflexivity and transitivity always reduce, they are not prevalues. And since coercion variables never reduce, they are prevalues. Finally, values are either prevalues or constructors applied to values. Notice that prevalues are neutral values, *i.e.* values that do not start with a constructor. We overload the notation from prevalues and values of the λ -calculus, but we make the distinction clear when necessary.

Reduction rules are given in Figure 4.3. We label the reduction with two annotations: β for computational steps and ι for typing steps. Computational steps are those of the λ -calculus,

¹see Section 6.1.7 for discussion about dependent type systems

$$\begin{array}{c}
\text{REDCTX} \\
\frac{M \rightsquigarrow_{\beta\iota} N}{E[M] \rightsquigarrow_{\beta\iota} E[N]}
\end{array}
\quad
\begin{array}{c}
\text{REDAPP} \\
(\lambda(x : \tau) M) N \rightsquigarrow_{\beta} M[x/N]
\end{array}
\quad
\begin{array}{c}
\text{REDFST} \\
\text{fst} \langle M, N \rangle \rightsquigarrow_{\beta} M
\end{array}
\quad
\begin{array}{c}
\text{REDSND} \\
\text{snd} \langle M, N \rangle \rightsquigarrow_{\beta} N
\end{array}$$

$$\begin{array}{c}
\text{REDREFL} \\
\langle M \rangle \rightsquigarrow_{\iota} M
\end{array}
\quad
\begin{array}{c}
\text{REDTRANS} \\
(G_2 \circ G_1) \langle M \rangle \rightsquigarrow_{\iota} G_2 \langle G_1 \langle M \rangle \rangle
\end{array}
\quad
\begin{array}{c}
\text{REDWEAK} \\
(*G) \langle M \rangle \rightsquigarrow_{\iota} G \langle M \rangle
\end{array}$$

Figure 4.3: Base system reduction rules

$$\begin{array}{c}
\text{TERMVAR} \\
\frac{\Gamma \text{env} \quad (x : \tau) \in \Gamma}{x \Rightarrow x : \Gamma \vdash \tau}
\end{array}
\quad
\begin{array}{c}
\text{TERMLAM} \\
\frac{M \Rightarrow a : \Gamma, (x : \tau) \vdash \sigma}{\lambda(x : \tau) M \Rightarrow \lambda x a : \Gamma \vdash \tau \rightarrow \sigma}
\end{array}$$

$$\begin{array}{c}
\text{TERMAPP} \\
\frac{M \Rightarrow a : \Gamma \vdash \tau \rightarrow \sigma \quad N \Rightarrow b : \Gamma \vdash \tau}{MN \Rightarrow ab : \Gamma \vdash \sigma}
\end{array}
\quad
\begin{array}{c}
\text{TERMPAIR} \\
\frac{M \Rightarrow a : \Gamma \vdash \tau \quad N \Rightarrow b : \Gamma \vdash \sigma}{\langle M, N \rangle \Rightarrow \langle a, b \rangle : \Gamma \vdash \tau \times \sigma}
\end{array}$$

$$\begin{array}{c}
\text{TERMFST} \\
\frac{M \Rightarrow a : \Gamma \vdash \tau \times \sigma}{\text{fst} M \Rightarrow \text{fst} a : \Gamma \vdash \tau}
\end{array}
\quad
\begin{array}{c}
\text{TERMSND} \\
\frac{M \Rightarrow a : \Gamma \vdash \tau \times \sigma}{\text{snd} M \Rightarrow \text{snd} a : \Gamma \vdash \sigma}
\end{array}$$

$$\begin{array}{c}
\text{TERMCOER} \\
\frac{M \Rightarrow a : \Gamma, \Sigma \vdash \tau \quad G \Rightarrow \Gamma \vdash (\Sigma \vdash \tau) \triangleright \sigma}{G \langle M \rangle \Rightarrow a : \Gamma \vdash \sigma}
\end{array}$$

Figure 4.4: Base system term judgment relation

while typing steps are new and only have to do with typings. We write $\beta\iota$ for either β or ι . The judgment for our reduction is thus written: $M \rightsquigarrow_{\beta\iota} M$.

Rule REDCTX is the context rule for strong reduction. If a term M can do a β (resp. ι) step to N , then the depth-one evaluation context E filled with M , namely $E[M]$, can also do a β (resp. ι) step to $E[N]$. Rules REDAPP, REDFST, and REDSND are similar to those of the λ -calculus, modulo the β annotation and the type annotation for the explicit abstraction.

Coercions never modify the computational content, and their reduction rules are thus ι -steps. Rule REDREFL shows that the reflexivity coercion does not modify a term and hence its typing. Rule REDTRANS shows that coercing the term M with $G_2 \circ G_1$ is like coercing M with G_1 and then coercing the result with G_2 , applying both coercions in sequence. Finally, rule REDWEAK simply removes the weakening coercion making the subcoercion accessible.

The term judgment relation, given in Figure 4.4, is of the form $M \Rightarrow a : \Gamma \vdash \tau$ to match the requested form $M \Rightarrow a : \Phi$ since invariants Φ are typings $\Gamma \vdash \tau$. The left part $M \Rightarrow$ is used for the explicit version only. The first six rules are those of the STLC modulo the different notation (see Section 3.1). The last rule, named TERMCOER, is new. It tells that a term of invariant $\Gamma, \Sigma \vdash \tau$ can be seen with invariant $\Gamma \vdash \sigma$, given a coercion from $\Gamma, \Sigma \vdash \tau$ to $\Gamma \vdash \sigma$; this is written $G \Rightarrow \Gamma \vdash (\Sigma \vdash \tau) \triangleright \sigma$.

The coercion judgment relation is written $G \Rightarrow \Gamma \vdash (\Sigma \vdash \tau) \triangleright \sigma$ and defined in Figure 4.5.

$$\begin{array}{c}
\text{COERREFL} \\
\Diamond \Rightarrow \Gamma \vdash \tau \triangleright \tau
\end{array}
\qquad
\frac{\text{COERTTRANS} \quad G_1 \Rightarrow \Gamma, \Sigma_2 \vdash (\Sigma_1 \vdash \tau_1) \triangleright \tau_2 \quad G_2 \Rightarrow \Gamma \vdash (\Sigma_2 \vdash \tau_2) \triangleright \tau_3}{G_2 \circ G_1 \Rightarrow \Gamma \vdash (\Sigma_2, \Sigma_1 \vdash \tau_1) \triangleright \tau_3}$$

$$\frac{\text{COERVAR} \quad (c : \tau \triangleright \sigma) \in \Gamma}{c \Rightarrow \Gamma \vdash \tau \triangleright \sigma}
\qquad
\frac{\text{COERWEAK} \quad \Gamma, \Sigma \text{ env} \quad G \Rightarrow \Gamma \vdash (\Sigma \vdash \tau) \triangleright \sigma}{*G \Rightarrow \Gamma \vdash \tau \triangleright \sigma}$$

Figure 4.5: Base system coercion judgment relation

When the erasable environment Σ is empty we may write $G \Rightarrow \Gamma \vdash \tau \triangleright \sigma$ instead of $G \Rightarrow \Gamma \vdash (\emptyset \vdash \tau) \triangleright \sigma$. The intuitive interpretation of $G \Rightarrow \Gamma \vdash (\Sigma \vdash \tau) \triangleright \sigma$ is that for all term a (and explicit term M) of typing $\Gamma, \Sigma \vdash \tau$, the term a (and explicit term $G\langle M \rangle$) has typing $\Gamma \vdash \sigma$. This interpretation comes from the fact that the typing rules of our considered extensions are of the form:

$$\frac{a : \Gamma, \Sigma \vdash \tau}{a : \Gamma \vdash \sigma}$$

For example, the typing rule for type abstraction is:

$$\frac{a : \Gamma, \alpha \vdash \tau}{a : \Gamma \vdash \forall \alpha \tau}$$

One may wonder why the coercion judgment is not written $\Gamma_1 \vdash \tau_1 \triangleright \Gamma_2 \vdash \tau_2$ as one could expect from our notion of coercion invariants $\Phi_1 \triangleright \Phi_2$. The reason is that we don't know which part of the environment is modified by the coercion, which is however necessary for weakening and substitution lemmas, for instance. For example, consider the coercion $(\Gamma, \alpha \vdash \tau) \triangleright (\Gamma, \alpha \vdash \forall \alpha \tau)$ which is sound, as we may generalize from $\Gamma, \alpha \vdash \tau$ to $\Gamma \vdash \forall \alpha \tau$ and weaken to $\Gamma, \alpha \vdash \forall \alpha \tau$. What if we weakened the coercion from Γ, α to $\Gamma, \alpha, (x : \alpha)$. It seems correct, but $(\Gamma, \alpha, (x : \alpha) \vdash \tau) \triangleright (\Gamma, \alpha, (x : \alpha) \vdash \forall \alpha \tau)$ is not. What we actually wanted to write was $\Gamma \vdash (\alpha \vdash \tau) \triangleright (\alpha \vdash \forall \alpha \tau)$ specifying that the environment Γ is not modified by the coercion. However, this notation $\Gamma \vdash (\Sigma_1 \vdash \tau_1) \triangleright (\Sigma_2 \vdash \tau_2)$ is subsumed by our current notation $\Gamma \vdash (\Sigma \vdash \tau) \triangleright \sigma$ using inverse modifications (see Section 6.1.11).

Rule COERREFL defines the reflexivity coercion \Diamond as the proof that typing $\Gamma \vdash \tau$ is included in itself without modifying the environment, which corresponds to the reflexivity closure of the inclusion relation. We remind that typings are to be interpreted as set of terms and coercions as inclusion proofs between these sets of terms.

Similarly, rule COERTTRANS defines the transitivity coercion $G_2 \circ G_1$ as the proof for inclusion transitivity. If G_1 proves that $\Gamma, \Sigma_2, \Sigma_1 \vdash \tau_1$ is included in $\Gamma, \Sigma_2 \vdash \tau_2$ by extending Σ_1 , and if G_2 proves that $\Gamma, \Sigma_2 \vdash \tau_2$ is included in $\Gamma \vdash \tau_3$ by extending Σ_2 , then $G_2 \circ G_1$ proves that $\Gamma, \Sigma_2, \Sigma_1 \vdash \tau_1$ is included in $\Gamma \vdash \tau_3$ by transitivity and extending Σ_2, Σ_1 .

Finally, rule COERVAR defines the coercion variable c as the proof that τ is included in σ , given that the environment Γ contains an hypothesis, named c , that τ is included in σ . And rule COERWEAK defines the weakening coercion $*G$ as the proof that τ is smaller than σ without any environment extension, whenever τ is smaller than σ with erasable environment

$$\begin{array}{c}
\text{TYPEVAR} \\
\frac{\Gamma \text{ env} \quad \alpha \in \Gamma}{\Gamma \vdash \alpha \text{ type}} \\
\\
\text{TYPEARR} \\
\frac{\Gamma \vdash \tau \text{ type} \quad \Gamma \vdash \sigma \text{ type}}{\Gamma \vdash \tau \rightarrow \sigma \text{ type}} \\
\\
\text{TYPEPROD} \\
\frac{\Gamma \vdash \tau \text{ type} \quad \Gamma \vdash \sigma \text{ type}}{\Gamma \vdash \tau \times \sigma \text{ type}} \\
\\
\text{ENVEMPTY} \\
\frac{}{\emptyset \text{ env}} \\
\\
\text{ENVTERM} \\
\frac{x \notin \text{dom}(\Gamma) \quad \Gamma \vdash \tau \text{ type}}{\Gamma, (x : \tau) \text{ env}}
\end{array}$$

Figure 4.6: Base system well-formedness relations

Σ and witness G . For the rule to be well-formed, the extended environment Γ, Σ has to be well-formed.

Type and environment well-formedness relations are given on Figure 4.6. We write $\Gamma \vdash \tau \text{ type}$ when type τ is well-formed under environment Γ . And environment well-formedness is written $\Gamma \text{ env}$. Type variable α is well-formed under environment Γ , by rule `TYPEVAR`, if it is bound in Γ and Γ is well-formed. Rules `TYPEARR` and `TYPEPROD` tell that $\tau \rightarrow \sigma$ and $\tau \times \sigma$ are well-formed under Γ if τ and σ are well-formed under Γ .

The empty environment is well-formed according to rule `ENVEMPTY`. Rule `ENVTERM` tells that the extended environment $\Gamma, (x : \tau)$ is well-formed if x is not already bound and τ is well-formed under Γ . Notice that there is an environment well-formedness rule for each kind of abstraction and not for each kind of variable.

A global note about the syntax of this base system is that features never *change* any kind of rules (typing rules, reduction rules, or well-formedness rules). Features may only *add* types (with their well-formedness rules), coercions (with their typing, ι -reduction, and environment rules), prevalues, and values. Finally they do not add more explicit terms, environments, erasable environments, or evaluation contexts. This non-modification property explains that the following features can be combined in an orthogonal way: all features rely on the base system only and no features alter the base system. The base system is sufficient for all extensions but contains only the expressivity of the STLC.

4.2 Features

Features usually define new invariants in the form of erasable types, like polymorphism which adds the polymorphic type $\forall \alpha \tau$. We make a distinction between *computational types* and *erasable types*. On the one hand, computational types (or simple types) are defined in the base system which corresponds to the STLC. We call them computational types, because their head constructor describe a computational behavior. The arrow type constructor classifies term abstractions, while the product type constructor classifies pairs of terms. Computational types are introduced and eliminated through term typing rules and each rule is associated to a term construct. These types and rules are already defined in the base system and no features will need a new one. Notice that we only consider type system features: a type system feature extends the type system but not the programming language. Adding integers to the programming language is not a type system extension, but a programming language extension. On the other hand, erasable types are related to typing and are independent from computation. They are introduced and eliminated through coercions and are thus erasable.

All features define an erasable type and its associated coercions. The only feature that does not add an erasable type is the η -expansion. This feature defines coercion rules for the congruence of computational types. As such, η -expansion could actually be considered in the base system. But since it adds subtyping, we decided to see it as an extension feature.

Congruence rules are actually only necessary for computational types. Erasable types have derivable congruence rules from their introduction and elimination rules. The typical case is using transitivity to eliminate the existing erasable constructor, then using the subcoercion, and finally introducing the erasable constructor back. In particular, congruence rules for polymorphic types and recursive types are derivable. For example, see Section 5.4.4 to see how to derive the congruence rule for recursive types.

4.2.1 Polymorphism

Polymorphism is the possibility to abstract over types in typing derivations. We make this new invariant apparent in the syntax with an erasable type. In the implicit version the only syntactical change is that we add the polymorphic type $\forall\alpha\tau$ to the syntax of types. For the explicit version, besides the polymorphic type, we need to add two coercions: one for type abstraction $\Lambda\alpha$ and one for type application $\cdot\tau$. We also need to extend prevalues and values accordingly.

$\tau, \sigma, \rho ::= \dots \mid \forall\alpha\tau$	Types
$G ::= \dots \mid \Lambda\alpha \mid \cdot\tau$	Coercions
$p ::= \dots \mid \cdot\tau\langle p \rangle$	Prevalues
$v ::= \dots \mid \Lambda\alpha\langle v \rangle$	Values

Reduction rules have to be extended with one additional rule called REDPOLY. If a type application $\cdot\tau$ follows a type abstraction $\Lambda\alpha$, then we can remove these nodes and substitute α for τ in the term.

$$\text{REDPOLY} \quad \cdot\tau\langle\Lambda\alpha\langle M \rangle\rangle \rightsquigarrow_{\iota} M[\alpha/\tau]$$

We add two coercion rules to the existing ones: one for type abstraction and one for type application. Rule COERTLAM is the most interesting one since it illustrates what there is to say about the coercion judgment. It can be read with rule TERMCOER under hand. If the term a has typing $\Gamma, \alpha \vdash \tau$ with witness M , then it also has typing $\Gamma \vdash \forall\alpha\tau$ with witness $\Lambda\alpha\langle M \rangle$. We see that both the type and the environment have been modified by the coercion. The environment has been extended with the type binding α , while the type gets closed by type abstraction. Similarly, rule COERTAPP tells that we can instantiate a polymorphic type $\forall\alpha\tau$ with a well-formed type σ . In this case, only the type is modified, the environment is not extended or modified in another way.

$\text{COERTLAM} \quad \Lambda\alpha \Rightarrow \Gamma \vdash (\alpha \vdash \tau) \triangleright \forall\alpha\tau$	$\text{COERTAPP} \quad \frac{\Gamma \vdash \sigma \text{ type}}{\cdot\sigma \Rightarrow \Gamma \vdash \forall\alpha\tau \triangleright \tau[\alpha/\sigma]}$
---	--

We also add rules to describe when a polymorphic type is well-formed and when an environment extended with a type variable is well-formed. A polymorphic type is well-formed if its body is well-formed under the environment extended with its type binding. An environment

extended with a type binding is well-formed if its prefix is well-formed and the type variable is not already bound.

$$\frac{\text{TYPEFOR}}{\Gamma, \alpha \vdash \tau \text{ type}} \quad \frac{\text{ENVTYPE} \quad \alpha \notin \text{dom}(\Gamma) \quad \Gamma \text{ env}}{\Gamma, \alpha \text{ env}}$$

4.2.2 Eta-expansion

We call η -expansion the extension of coercions with computational type congruence. The reason why we use this name is because η -expansion derivations give the congruence rules of computational types. This feature is actually more expressive than the usual congruence rules we can find in type systems with subtyping. Because in our setting, coercions are not only between types, but between typings (the pair of an environment and a type). So η -expansion allows us to do subtyping between typings and not only types. We will explain in a few paragraph how η -expansion works.

This extension does not need any syntactic changes in the implicit version. However, we need to add two coercions in the explicit version. Actually, we need to add as many coercions as we have computational types. In our case, we have just two computational types: the arrow type and the product type. So we have two η -expansion coercions, one for each.

$$G ::= \dots \mid G \xrightarrow{\tau} G \mid G \times G \quad \text{Coercions}$$

These coercions can be intuitively understood as $\lambda(x : \tau) G_2 \langle \langle \rangle (G_1 \langle x \rangle) \rangle$ for the arrow η -expansion coercion $G_1 \xrightarrow{\tau} G_2$, and $\langle G_1 \langle \text{fst } \langle \rangle \rangle, G_2 \langle \text{snd } \langle \rangle \rangle \rangle$ for the product η -expansion coercion $G_1 \times G_2$. These are called η -expansion because when we take their image by the drop function erasing types and coercions, we get the η -expansions $\lambda x \langle \rangle x$ for the arrow type and $\langle \text{fst } \langle \rangle, \text{snd } \langle \rangle \rangle$ for the product type. This gives an idea of why these coercions are erasable and do not modify the computational content of their hole.

This intuition could be used as an encoding: whenever we want to retype a function, we can η -expand it. However, the resulting term would be an η -expansion of the initial term, and this disagrees with our wish of coercions being truly erasable. Moreover, when coercing (or subtyping) deeply in a type, it is more intuitive to have a notation close to the structure of the type, than to write the associated η -expansion.

Before we look at the associated coercion rules, we need to extend the prevalues. Several choices can be done depending on the reduction rules we want and this is discussed in Section 4.3 when describing the reductions rules of F_t^P . In our case, η -expansions are only destructors, so they are prevalues.

$$p ::= \dots \mid (G \xrightarrow{\tau} G) \langle p \rangle \mid (G \times G) \langle p \rangle \quad \text{Prevalues}$$

Since the η -expansion coercions are destructors, they reduce when they are applied to their associated constructor in accordance with their intuitive definition. In other words, we take their η -expansion view as terms and fill the hole with their argument (starting with the correct constructor) and reduce the created redex at the hole. For instance, for rule `REDETAARR`, the left-hand side $(G_1 \xrightarrow{\tau'} G_2) \langle \lambda(x : \tau) M \rangle$ can be seen as $\lambda(x : \tau') G_2 \langle (\lambda(x : \tau) M) (G_1 \langle x \rangle) \rangle$ which

reduces to the right-hand side $\lambda(x : \tau') G_2 \langle M[x/G_1 \langle x \rangle] \rangle$.

$$\text{REDETAARR} \\ (G_1 \xrightarrow{\tau'} G_2) \langle \lambda(x : \tau) M \rangle \rightsquigarrow_\iota \lambda(x : \tau') G_2 \langle M[x/G_1 \langle x \rangle] \rangle$$

$$\text{REDETAProd} \\ (G_1 \times G_2) \langle \langle M_1, M_2 \rangle \rangle \rightsquigarrow_\iota \langle G_1 \langle M_1 \rangle, G_2 \langle M_2 \rangle \rangle$$

When adding a new feature, the most important part of the definition is always the extension of the coercion relation. Here, the rules follow the intuition of the η -expansion view. Rule COERETAARR tells that an arrow type $\tau \rightarrow \sigma$ under environment Γ, Σ can be coerced into the arrow type $\tau' \rightarrow \sigma'$ under the environment Γ , given that τ' can be coerced to τ under the extended environment Γ, Σ and that σ can be coerced to σ' under Γ and binding the erasable environment Σ . Notice the contravariance for the left-hand type of the arrow constructor. For well-formedness reasons, the type τ' has to be well-formed under environment Γ in order for the type variables bound in Σ to not be used by τ' . One may wonder why G_2 binds Σ in G_1 and why G_1 does not bind anything. This can be understood by looking at the erasable context view $\lambda(x : \tau) G_2 \langle [] (G_1 \langle x \rangle) \rangle$. We see that G_1 is actually under the scope of G_2 , and that the only thing under the scope of G_1 is the term variable x , which does not use anything from its environment but itself.

Rule COERETAPROD is easier since the product type is covariant on both arguments. If the typing $\Gamma, \Sigma \vdash \tau$ can be coerced to the typing $\Gamma \vdash \tau'$ and similarly the typing $\Gamma, \Sigma \vdash \sigma$ can be coerced to $\Gamma \vdash \sigma'$, then the product type $\tau \times \sigma$ can be coerced to the product $\tau' \times \sigma'$ under the environment Γ and extending with the erasable environment Σ .

$$\text{COERETAARR} \\ \frac{\Gamma \vdash \tau' \text{ type} \quad G_1 \Rightarrow \Gamma, \Sigma \vdash (\emptyset \vdash \tau') \triangleright \tau \quad G_2 \Rightarrow \Gamma \vdash (\Sigma \vdash \sigma) \triangleright \sigma'}{G_1 \xrightarrow{\tau'} G_2 \Rightarrow \Gamma \vdash (\Sigma \vdash \tau \rightarrow \sigma) \triangleright \tau' \rightarrow \sigma'}$$

$$\text{COERETAPROD} \\ \frac{G_1 \Rightarrow \Gamma \vdash (\Sigma \vdash \tau) \triangleright \tau' \quad G_2 \Rightarrow \Gamma \vdash (\Sigma \vdash \sigma) \triangleright \sigma'}{G_1 \times G_2 \Rightarrow \Gamma \vdash (\Sigma \vdash \tau \times \sigma) \triangleright \tau' \times \sigma'}$$

Notice that when adding η -expansion and polymorphism to the base system, we get exactly System F_η . This will be shown in Section 4.5.2. However, this presentation differs from the original presentation in [25] by factorizing free generalization, type instantiation, arrow congruence, distributivity, and polymorphic congruence into type generalization, type instantiation, and arrow η -expansion. This presentation is more economical, which has been made possible by extending coercions from types to typings.

4.2.3 Bottom

This extension closes the hierarchy of types with a minimum, called bottom and written \perp . This extension is useful later when dealing with lower bounded polymorphism. There is also a similar extension to close the hierarchy of types with a maximum, called top and written \top (see the next section). We extend the syntax of types with the bottom type \perp . We extend coercions with the absurd coercion $\perp \tau$ and prevalues with $(\perp \tau) \langle p \rangle$.

$\tau, \sigma, \rho ::= \dots \mid \perp$	Types
$G ::= \dots \mid \perp\tau$	Coercions
$p ::= \dots \mid (\perp\tau)\langle p \rangle$	Prevalues

There are no reduction rules associated to the bottom type because there is no constructor of type bottom. But there is a coercion rule `CoERBot` telling that from the bottom type we can go to any type. So \perp is actually a minimum of all well-formed types in any environment, according to the coercion relation we are defining.

$$\frac{\text{CoERBot} \quad \Gamma \vdash \tau \text{ type}}{\perp\tau \Rightarrow \Gamma \vdash \perp \triangleright \tau}$$

Finally, we need to give the well-formedness rule for the bottom type, namely `TypeBot`. The bottom type is well-formed under any well-formed environment.

$$\frac{\text{TypeBot} \quad \Gamma \text{ env}}{\Gamma \vdash \perp \text{ type}}$$

4.2.4 Top

This extension is similar to the preceding one and closes the hierarchy of types with a maximum, called top and written \top . This extension is useful later, and in particular when dealing with upper bounded polymorphism. We extend the syntax of types with the top type \top . We extend coercions with the forget coercion \top and values with the blind value $\top\langle v \rangle$.

$\tau, \sigma, \rho ::= \dots \mid \top$	Types
$G ::= \dots \mid \top$	Coercions
$v ::= \dots \mid \top\langle v \rangle$	Values

There are no reduction rules associated to the top type because there are no associated destructors. But there is a coercion rule `CoERTop` that tells that any type can be coerced to top. So \top is actually a maximum of all well-formed types in any environment.

$$\frac{\text{CoERTop}}{\top \Rightarrow \Gamma \vdash \tau \triangleright \top}$$

Finally, we need to give the well-formedness rule for the top type, namely `TypeTop`. The top type is well-formed under any well-formed environment.

$$\frac{\text{TypeTop} \quad \Gamma \text{ env}}{\Gamma \vdash \top \text{ type}}$$

4.2.5 Lower Bounded polymorphism

Lower bounded polymorphism, when paired with the bottom type on top of the base system, corresponds to MLF. We show in detail how in Section 4.5.3. A small difference lies in the absence of recursive bounds in MLF which are possible in this extension.

This extension resembles the extension of the STLC with polymorphism but instead of abstracting over a type, we abstract over a type greater than another one. In other words, we abstract over instances of a type, called the lower bound. We say that the bound is recursive when it may contain occurrences of the abstract type. We extend the syntax of types with the lower bounded polymorphic type $\forall(\alpha \triangleleft \tau) \tau$. We add two coercions: one for abstraction $\Lambda(\alpha \triangleleft c : \tau)$ and one for application $\cdot[\tau \triangleleft G]$. Notice that in $\Lambda(\alpha \triangleleft c : \tau)$ we simultaneously abstract over c and α . We extend the prevalues and values accordingly.

$\tau, \sigma, \rho ::= \dots \mid \forall(\alpha \triangleleft \tau) \tau$	Types
$G ::= \dots \mid \Lambda(\alpha \triangleleft c : \tau) \mid \cdot[\tau \triangleleft G]$	Coercions
$p ::= \dots \mid (\cdot[\tau \triangleleft G])\langle p \rangle$	Prevalues
$v ::= \dots \mid \Lambda(\alpha \triangleleft c : \tau)\langle v \rangle$	Values

We extend the reduction rules with rule REDPOLYL which is similar to REDPOLY; it takes a lower bounded type abstraction followed by a lower bounded type application, and it reduces it into a type substitution followed by a coercion substitution.

$$\text{REDPOLYL} \quad \cdot[\sigma \triangleleft G]\langle \Lambda(\alpha \triangleleft c : \tau)\langle M \rangle \rangle \rightsquigarrow_{\iota} M[\alpha/\sigma][c/G]$$

The coercion typing rules need also to be extended. Rule COERTLAML tells that if we can type a term with ρ under an environment extended with type variable α and coercion variable c of coercion type $\tau \triangleright \alpha$, then this term also has type $\forall(\alpha \triangleleft \tau) \rho$ under the non-extended environment. Notice that this abstraction extends the environment with two binders and, although we write $\alpha \triangleleft \tau$ in the syntax of types, we still write $\tau \triangleright \alpha$ in the coercion binding. Rule COERTAPPL permits to view a term of type $\forall(\alpha \triangleleft \tau) \rho$ with type $\rho[\alpha/\sigma]$ whenever σ is well-formed and an instance of $\tau[\alpha/\sigma]$. Here, we must substitute α by σ in τ since we allow recursive bounds.

$\text{COERTLAML} \quad \Lambda(\alpha \triangleleft c : \tau) \Rightarrow \Gamma \vdash (\alpha, (c : \tau \triangleright \alpha) \vdash \rho) \triangleright \forall(\alpha \triangleleft \tau) \rho$	$\text{COERTAPPL} \quad \frac{\Gamma \vdash \sigma \text{ type} \quad G \Rightarrow \Gamma \vdash \tau[\alpha/\sigma] \triangleright \sigma}{\cdot[\sigma \triangleleft G] \Rightarrow \Gamma \vdash \forall(\alpha \triangleleft \tau) \rho \triangleright \rho[\alpha/\sigma]}$
---	---

Finally, we extend the well-formedness rules. By rule TYPEFORL, the lower bounded polymorphic type $\forall(\alpha \triangleleft \tau) \rho$ is well-formed if its body ρ is well-formed under the environment extended with the type binding α and coercion binding $(c : \tau \triangleright \alpha)$. So α is bound in ρ and we know that it is an instance of τ . By rule ENVTYPEL, the extended environment $\Gamma, \alpha, (c : \tau \triangleright \alpha)$ is well-formed if the bound variables are not already bound in environment Γ and τ is well-formed under Γ, α which allows τ to mention α and thus permits recursive bounds.

$\text{TYPEFORL} \quad \frac{\Gamma, \alpha, (c : \tau \triangleright \alpha) \vdash \rho \text{ type}}{\Gamma \vdash \forall(\alpha \triangleleft \tau) \rho \text{ type}}$	$\text{ENVTYPEL} \quad \frac{c, \alpha \notin \text{dom}(\Gamma) \quad \Gamma, \alpha \vdash \tau \text{ type}}{\Gamma, \alpha, (c : \tau \triangleright \alpha) \text{ env}}$
--	--

4.2.6 Upper Bounded polymorphism

Upper bounded polymorphism, when paired with top type and η -expansion on top of the base system, gives a type system more expressive than the most expressive version of System $F_{<}$. We show in detail how in Section 4.5.4.

This extension resembles lower bounded polymorphism but instead of abstracting over a type greater than another one, it abstracts over a type smaller than another one. We call this last type the upper bound of the abstract type. We say that the bound is recursive when it may contain occurrences of the abstract type. We extend the syntax of types with the upper bounded polymorphic type $\forall(\alpha \triangleright \tau) \tau$. We add two coercions: one for abstraction $\Lambda(\alpha \triangleright c : \tau)$ and one for application $\cdot[\tau \triangleright G]$. We also extend the prevalues and values accordingly and similarly to the previous extension. Notice that only the triangles are inverted from \triangleleft to \triangleright in a way to make what were previously lower bounds upper bounds.

$\tau, \sigma, \rho ::= \dots \mid \forall(\alpha \triangleright \tau) \tau$	Types
$G ::= \dots \mid \Lambda(\alpha \triangleright c : \tau) \mid \cdot[\tau \triangleright G]$	Coercions
$p ::= \dots \mid (\cdot[\tau \triangleright G])\langle p \rangle$	Prevalues
$v ::= \dots \mid \Lambda(\alpha \triangleright c : \tau)\langle v \rangle$	Values

We extend the reduction rules with rule REDPOLYU which is similar to REDPOLYL since it reduces an upper bounded type abstraction followed by an upper bounded type application into a type substitution followed by a coercion substitution.

$$\text{REDPOLYU} \quad \cdot[\sigma \triangleright G]\langle \Lambda(\alpha \triangleright c : \tau)\langle M \rangle \rangle \rightsquigarrow_{\iota} M[\alpha/\sigma][c/G]$$

The coercion typing rules need also to be extended. Rule COERTLAMU says that if we can type a term with ρ under an environment extended with type variable α and coercion variable c of coercion type $\alpha \triangleright \tau$, then this term also has type $\forall(\alpha \triangleright \tau) \rho$ under the non-extended environment. Notice that this abstraction extend the environment with two binders similarly to rule COERTLAML. Rule COERTAPPU permits to view a term of type $\forall(\alpha \triangleright \tau) \rho$ with type $\rho[\alpha/\sigma]$ whenever σ is well-formed and smaller than $\tau[\alpha/\sigma]$. We substituted α by σ in τ since we allow recursive bounds.

$$\begin{array}{c} \text{COERTLAMU} \\ \Lambda(\alpha \triangleright c : \tau) \Rightarrow \Gamma \vdash (\alpha, (c : \alpha \triangleright \tau) \vdash \rho) \triangleright \forall(\alpha \triangleright \tau) \rho \end{array} \qquad \begin{array}{c} \text{COERTAPPU} \\ \Gamma \vdash \sigma \text{ type} \quad G \Rightarrow \Gamma \vdash \sigma \triangleright \tau[\alpha/\sigma] \\ \hline \cdot[\sigma \triangleright G] \Rightarrow \Gamma \vdash \forall(\alpha \triangleright \tau) \rho \triangleright \rho[\alpha/\sigma] \end{array}$$

Finally, we extend the well-formedness rules. By rule TYPEFORU, the upper bounded polymorphic type $\forall(\alpha \triangleright \tau) \rho$ is well-formed if its body ρ is well-formed under the environment extended with the type binding α and coercion binding $(c : \alpha \triangleright \tau)$. So α is bound in ρ and we know that it is smaller than τ . By rule ENVTYPEU, the extended environment $\Gamma, \alpha, (c : \alpha \triangleright \tau)$ is well-formed if the bound variables are not already bound in environment Γ and τ is well-formed under Γ, α which allows τ to mention α and thus permits recursive bounds.

$$\begin{array}{c} \text{TYPEFORU} \\ \Gamma, \alpha, (c : \alpha \triangleright \tau) \vdash \rho \text{ type} \\ \hline \Gamma \vdash \forall(\alpha \triangleright \tau) \rho \text{ type} \end{array} \qquad \begin{array}{c} \text{ENVTYPEU} \\ c, \alpha \notin \text{dom}(\Gamma) \quad \Gamma, \alpha \vdash \tau \text{ type} \\ \hline \Gamma, \alpha, (c : \alpha \triangleright \tau) \text{ env} \end{array}$$

4.3 System F_{ι}^p

Now that the base system and the features are presented, we may compose them altogether into a type system that we call System F_{ι}^p . This language is described with a slightly different

α, β	Type variables
c	Coercion variables
$M, N ::= x \mid \lambda(x : \tau) M \mid M M \mid \langle M, M \rangle \mid \text{fst } M \mid \text{snd } M \mid G \langle M \rangle$	Explicit terms
$\tau, \sigma, \rho ::= \alpha \mid \tau \rightarrow \tau \mid \tau \times \tau \mid \forall \alpha \tau \mid \perp \mid \top \mid \forall(\alpha \triangleleft \tau) \tau \mid \forall(\alpha \triangleright \tau) \tau$	Types
$G ::= c \mid \diamond \mid G \circ G \mid *G \mid \Lambda \alpha \mid \cdot \tau \mid G \xrightarrow{\tau} G \mid G \times G \mid \perp \tau \mid \top$ $\mid \Lambda(\alpha \triangleleft c : \tau) \mid \cdot[\tau \triangleleft G] \mid \Lambda(\alpha \triangleright c : \tau) \mid \cdot[\tau \triangleright G]$	Coercions
$\Gamma ::= \emptyset \mid \Gamma, (x : \tau) \mid \Gamma, \alpha \mid \Gamma, (c : \tau \triangleright \tau)$	Environments

Figure 4.7: System \mathbb{F}_l^P syntax

presentation in a paper [13]. In this framework, we made it clear that coercions act on typings and not only on types.

We sum up the syntax of the base system and all extensions in Figure 4.7. We write type variables α or β and coercion variables c . Coercion variables are only necessary in the explicit version of the type system. Terms are written M or N and only necessary in the explicit version. They contain variables x , abstractions $\lambda(x : \tau) M$, applications $M M$, pairs $\langle M, M \rangle$, projections $\text{fst } M$ and $\text{snd } M$, and coercion constructs $G \langle M \rangle$.

Types are written τ , σ , or ρ . They contain variables α , arrow types $\tau \rightarrow \tau$, product types $\tau \times \tau$, polymorphic types $\forall \alpha \tau$, the bottom type \perp , the top type \top , lower bounded polymorphic types $\forall(\alpha \triangleright \tau) \tau$, and upper bounded polymorphic types $\forall(\alpha \triangleleft \tau) \tau$. In bounded polymorphic types $\forall(\alpha \triangleright \tau) \rho$ and $\forall(\alpha \triangleleft \tau) \rho$, the type variable α is bound in both the bound τ and the body ρ . It is bound in ρ because it is a polymorphic type, and it is bound in τ because we allow recursive bounds. The upper (resp. lower) bound polymorphic type $\forall(\alpha \triangleright \tau) \rho$ (resp. $\forall(\alpha \triangleleft \tau) \rho$) can be read as type ρ for all abstract type α such that a coercion from α to τ (resp. from τ to α) exists.

Coercions are written G and only necessary in the explicit version since they only occur in the coercion constructs of terms which are by definition only necessary in the explicit version. Coercions contain variables c , the reflexivity coercion \diamond , transitivity coercions $G \circ G$, weakenings $*G$, type abstractions $\Lambda \alpha$, type applications $\cdot \tau$, arrow η -expansions $G \xrightarrow{\tau} G$, product η -expansions $G \times G$, bottom coercions $\perp \tau$, the top coercion \top , lower bounded type abstractions $\Lambda(\alpha \triangleleft c : \tau)$, lower bounded type applications $\cdot[\tau \triangleleft G]$, upper bounded type abstractions $\Lambda(\alpha \triangleright c : \tau)$, and upper bounded type applications $\cdot[\tau \triangleright G]$.

The transitivity coercions have to be read from right to left. The coercion $G_2 \circ G_1$ means that G_2 occurs after G_1 , as we can see in rule REDTRANS. The type annotation on the arrow η -expansion coercion is needed for rule REDARR. Upper and lower bounded abstractions bind the type variable α and the coercion variable c , which has type $\alpha \triangleright \tau$ and $\tau \triangleright \alpha$ respectively. Notice that for lower bounded abstraction, we write $\Lambda(\alpha \triangleleft c : \tau)$ with a reverse \triangleright . This is because we want to enhance that both c and α are bound and that c has coercion type $\tau \triangleright \alpha$. A similar reason holds for the lower bounded polymorphic type $\forall(\alpha \triangleleft \tau) \rho$. Upper and lower bounded applications are made of three parts: a type instantiation τ , a coercion instantiation G , and a direction \triangleright or \triangleleft which corresponds to upper or lower bounded application, respectively. The orientation is the same as its associated abstraction.

Finally, we define environments Γ as lists of bindings. The empty environment is written \emptyset and extended environments are written with a comma. Environments extended with a term binding are written $\Gamma, (x : \tau)$, environments extended with a type binding are written Γ, α ,

$\Sigma ::= \emptyset \mid \Sigma, \alpha \mid \Sigma, (c : \tau \triangleright \tau)$	Erasable environments
$E ::= \lambda(x : \tau) \square \mid \square M \mid M \square \mid \langle \square, M \rangle \mid \langle M, \square \rangle \mid \mathbf{fst} \square \mid \mathbf{snd} \square \mid G \langle \square \rangle$	Evaluation contexts
$p ::= x \mid p v \mid \mathbf{fst} p \mid \mathbf{snd} p \mid \cdot \tau \langle p \rangle \mid (G \xrightarrow{\tau} G) \langle p \rangle \mid (G \times G) \langle p \rangle$ $\quad \mid (\perp \tau) \langle p \rangle \mid (\cdot [\tau \triangleleft G]) \langle p \rangle \mid (\cdot [\tau \triangleright G]) \langle p \rangle \mid c \langle v \rangle$	Prevalues
$v ::= p \mid \lambda(x : \tau) v \mid \langle v, v \rangle \mid \Lambda \alpha \langle v \rangle \mid \top \langle v \rangle \mid \Lambda(\alpha \triangleleft c : \tau) \langle v \rangle \mid \Lambda(\alpha \triangleright c : \tau) \langle v \rangle$	Values

Figure 4.8: System F_t^p notations

RED_{CTX} $\frac{M \rightsquigarrow_{\beta_t} N}{E[M] \rightsquigarrow_{\beta_t} E[N]}$	RED_{APP} $(\lambda(x : \tau) M) N \rightsquigarrow_{\beta} M[x/N]$	RED_{FST} $\mathbf{fst} \langle M, N \rangle \rightsquigarrow_{\beta} M$	RED_{SND} $\mathbf{snd} \langle M, N \rangle \rightsquigarrow_{\beta} N$
RED_{REFL} $\diamond \langle M \rangle \rightsquigarrow_t M$	$\text{RED}_{\text{TRANS}}$ $(G_2 \circ G_1) \langle M \rangle \rightsquigarrow_t G_2 \langle G_1 \langle M \rangle \rangle$	RED_{WEAK} $(*G) \langle M \rangle \rightsquigarrow_t G \langle M \rangle$	
RED_{POLY} $\cdot \tau \langle \Lambda \alpha \langle M \rangle \rangle \rightsquigarrow_t M[\alpha/\tau]$	$\text{RED}_{\text{ETAARR}}$ $(G_1 \xrightarrow{\tau'} G_2) \langle \lambda(x : \tau) M \rangle \rightsquigarrow_t \lambda(x : \tau') G_2 \langle M[x/G_1 \langle x \rangle] \rangle$		
$\text{RED}_{\text{ETAPROD}}$ $(G_1 \times G_2) \langle \langle M_1, M_2 \rangle \rangle \rightsquigarrow_t \langle G_1 \langle M_1 \rangle, G_2 \langle M_2 \rangle \rangle$	$\text{RED}_{\text{POLYL}}$ $\cdot [\sigma \triangleleft G] \langle \Lambda(\alpha \triangleleft c : \tau) \langle M \rangle \rangle \rightsquigarrow_t M[\alpha/\sigma][c/G]$		
	$\text{RED}_{\text{POLYU}}$ $\cdot [\sigma \triangleright G] \langle \Lambda(\alpha \triangleright c : \tau) \langle M \rangle \rangle \rightsquigarrow_t M[\alpha/\sigma][c/G]$		

Figure 4.9: System F_t^p reduction rules

and environments extended with a coercion binding are written $\Gamma, (c : \tau \triangleright \tau)$ in the explicit version and $\Gamma, (\tau \triangleright \tau)$ in the implicit version. Term bindings $(x : \tau)$ bind the term variable x to its type τ . A type binding α binds the type variable α . And coercion bindings $(c : \tau \triangleright \sigma)$ bind the coercion variable c to the coercion hypothesis $\tau \triangleright \sigma$. All variables bound in an environment have to be distinct. We ensure this restriction in the well-formedness judgment of environments.

We define some notations in Figure 4.8. Erasable environments are environments containing only erasable bindings. Erasable bindings are type or coercion bindings since they have to do with typing. Non-erasable bindings are term bindings since they have to do with computation. We write Σ for erasable environments.

We define evaluation contexts, prevalues, and values. All of these are defined according to strong reduction. In strong reduction, evaluation contexts are all one-hole contexts of depth one. Prevalues are destructors applied to prevalues where a constructor is expected and values elsewhere. Finally, values are constructors applied to values. For System F_t^p , prevalues contain in particular prevalues applied to a type $\cdot \tau \langle p \rangle$, prevalues applied to an η -expansion $(G \xrightarrow{\tau} G) \langle p \rangle$ and $(G \times G) \langle p \rangle$, instantiations of absurd prevalues $(\perp \tau) \langle p \rangle$, bounded type applications of prevalues $(\cdot [\tau \triangleright G]) \langle p \rangle$ and $(\cdot [\tau \triangleleft G]) \langle p \rangle$, and coerced values with an abstract coercion $c \langle v \rangle$. And values contain in particular type abstractions of values $\Lambda \alpha \langle v \rangle$, top values $\top \langle v \rangle$, and bounded type abstractions of values $\Lambda(\alpha \triangleright c : \tau) \langle v \rangle$ and $\Lambda(\alpha \triangleleft c : \tau) \langle v \rangle$.

We define the reduction rules in Figure 4.9. They are labeled with a β annotation for

computational steps and ι annotation for typing steps or erasable steps. We use the meta-variable $\beta\iota$ to designate either of these annotations.

The first four rules mimic those of the λ -calculus. Rule REDCTX is the context rule and simply transfers the annotation from its subreduction to its whole reduction. If M reduces to N with annotation β (resp. ι), then $E[M]$ reduces to $E[N]$ with annotation β (resp. ι). Rules REDAPP, REDFST, and REDSND are β -reduction rules since they actually do a computational step. All following reduction rules are ι -reduction rules and only have to do with typings.

Rules REDREFL and REDTRANS have to do with the closure properties of the coercion relation. The reflexivity coercion closes the coercion relation by reflexivity: typing $\Gamma \vdash \tau$ is smaller than typing $\Gamma \vdash \tau$ by coercion proof \diamond . While the transitivity coercion closes the coercion relation by transitivity: if typing $\Gamma_1 \vdash \tau_1$ is smaller than typing $\Gamma_2 \vdash \tau_2$ by coercion proof G_1 and typing $\Gamma_2 \vdash \tau_2$ is smaller than typing $\Gamma_3 \vdash \tau_3$ by coercion proof G_2 , then typing $\Gamma_1 \vdash \tau_1$ is smaller than typing $\Gamma_3 \vdash \tau_3$ by coercion proof $G_2 \circ G_1$. This is why $\diamond\langle M \rangle$ ι -reduces to M since the typing does not change, and why $(G_2 \circ G_1)\langle M \rangle$ ι -reduces to $G_2\langle G_1\langle M \rangle \rangle$. Rule REDWEAK simply applies the subcoercion.

Rule REDPOLY corresponds to the usual rule of System F for polymorphism. A type abstraction followed by a type application results in a type substitution. In terms of typings and coercions, if we change the typing of a term M from typing $\Gamma, \alpha \vdash \tau$ to typing $\Gamma \vdash \forall \alpha \tau$ by coercion proof $\Lambda \alpha$ and then from typing $\Gamma \vdash \forall \alpha \tau$ to typing $\Gamma \vdash \tau[\alpha/\sigma]$ by coercion proof $\cdot \sigma$ where σ is a well-formed type under environment Γ , then the substitution of the free occurrences of the type variable α by the type σ , written $[\alpha/\sigma]$, changes the typing $\Gamma, \alpha \vdash \tau$ of the term M to the typing $\Gamma \vdash \tau[\alpha/\sigma]$, which is the same as what the original coercions were doing.

Rules REDETAARR and REDETAProd deal with η -expansion. They can be read using the η -expansion intuition. Coercions are erasable contexts and the η -expansion coercion $G_1 \xrightarrow{\tau'} G_2$ can be seen as the erasable context $\lambda(x : \tau') G_2\langle \square \rangle (G_1\langle x \rangle)$. When we fill this context with $\lambda(x : \tau) M$ as it is the case in rule REDETAARR we get $\lambda(x : \tau') G_2\langle (\lambda(x : \tau) M) \rangle (G_1\langle x \rangle)$ which reduces to $\lambda(x : \tau') G_2\langle M[x/G_1\langle x \rangle] \rangle$ which is exactly the right-hand side of the reduction rule. The same mechanism works for rule REDETAProd when taking $\langle G_1\langle \text{fst } \square \rangle, G_2\langle \text{snd } \square \rangle \rangle$ as the erasable context for coercion $G_1 \times G_2$.

Some additional or alternate rules could have been used instead of rules REDETAARR and REDETAProd. We focus on the arrow type only, since the product type is very similar. The two additional or alternate rules we consider are REDETAARR1 and REDETAARR2.

$$\begin{array}{c} \text{REDETAARR1} \\ (G_1 \xrightarrow{\tau'} G_2)\langle M \rangle N \rightsquigarrow_{\iota} G_2\langle M (G_1\langle N \rangle) \rangle \end{array}$$

Rule REDETAARR1 is the analog of rule REDETAARR in the sense that an η -expansion has two potential redexes: one with the destructor above its hole, one with the constructor at its root. Rule REDETAARR considers the redex under its hole while rule REDETAARR1 considers the redex at its root. If we write down the η -expansion view of rule REDETAARR1 as erasable contexts we get $(\lambda(x : \tau') G_2\langle M (G_1\langle x \rangle) \rangle) N$ which reduces at its root to $G_2\langle M (G_1\langle N \rangle) \rangle$. Notice that one advantage of this rule on rule REDETAARR is that the type annotation τ' is not needed and the only reason for this annotation to be present in the syntax is to formulate rule REDETAARR. However, rule REDETAARR1 modifies the set of prevalues and values and in particular the classification of irreducible terms according to their type. For instance $(G_1 \xrightarrow{\tau'} G_2)\langle \lambda(x : \tau) M \rangle$ is an irreducible well-typed term if we replace rule REDETAARR with

rule REDETAARR1. Both rules, REDETAARR and REDETAARR1, can be used independently or together.

$$\text{REDETAARR2} \\ (G'_1 \xrightarrow{\tau'} G'_2) \langle (G_1 \xrightarrow{\tau} G_2) \langle M \rangle \rangle \rightsquigarrow_{\iota} ((G_1 \circ G'_1) \xrightarrow{\tau'} (G'_2 \circ G_2)) \langle M \rangle$$

Rule REDETAARR2 is special and actually problematic. Rule REDETAARR1 was the analog of rule REDETAARR because it was using the other redex of the hidden η -expansion, whereas rule REDETAARR2 mixes the root half-redex of one η -expansion with the hole half-redex of another η -expansion. If we write down the η -expansion view of this rule as erasable contexts we get:

$$\lambda(x : \tau') G'_2 \langle (\lambda(x : \tau) G_2 \langle M (G_1 \langle x \rangle) \rangle) (G'_1 \langle x \rangle) \rangle$$

It reduces to $\lambda(x : \tau') G'_2 \langle G_2 \langle M (G_1 \langle G'_1 \langle x \rangle \rangle) \rangle \rangle$ which is a reduct of the right-hand side of the reduction rule. This is where there is a problem since we break our intuition: coercions are erasable contexts. One way to recover it is to ask the closure reductions to be actually closure equivalences. In particular, we would have the following equivalence rules.

$$\diamond \langle M \rangle \equiv M \quad (G_2 \circ G_1) \langle M \rangle \equiv G_2 \langle G_1 \langle M \rangle \rangle \quad G_3 \circ (G_2 \circ G_1) \equiv (G_3 \circ G_2) \circ G_1$$

Rule REDETAARR2 would then be compatible with both REDETAARR and REDETAARR1. However, this rule alone is not enough to ensure progress since it does not act on computational terms as an η -expansion rule should do. So the six possible configurations for the η -expansion rules of the arrow type are: rule REDETAARR alone, rule REDETAARR1 alone, rules REDETAARR and REDETAARR1 together, rule REDETAARR2 with rule REDETAARR, rule REDETAARR2 with rule REDETAARR1, and rule REDETAARR2 with rules REDETAARR and REDETAARR1. And each time we add rule REDETAARR2 we have to modify the framework to allow term equivalence and add enough structural equivalence rules. Also, for each situation, the set of prevalues and values differ.

Rules REDPOLYU and REDPOLYL are similar to rule REDPOLY. A bounded type abstraction followed by a similarly bounded type application results in a type substitution followed by a coercion substitution.

Term typing rules are given on Figure 4.10. The term typing judgment is of the form $M \Rightarrow a : \Gamma \vdash \tau$ where M is a term which is a partial proof that the term a has type τ under environment Γ . It is a partial proof because it only contains enough information to rebuild the term a , and enough information to rebuild the type τ given an environment Γ . The term is necessary only in the explicit version of the type system.

Rule TERMVAR gives type τ to the term variable x if x is bound to τ in environment Γ . The term witnessing this rule is the term variable x itself. Since we want derivations of the judgment $M \Rightarrow a : \Gamma \vdash \tau$ to hold the proof that its type τ is well-formed under its environment Γ , we have to ask the environment Γ to be well-formed. Rule TERMLAM gives type $\tau \rightarrow \sigma$ to the term $\lambda x a$ under environment Γ if the term a has type σ under the extended environment $\Gamma, (x : \tau)$ where the term variable x is now bound to the type τ . If we have M the term for the premise, we write $\lambda(x : \tau) M$ the term witnessing the conclusion. We need to add the type annotation τ since we cannot rebuild it from M . Rule TERMAPP gives type σ to the application of a to b under environment Γ if a has type $\tau \rightarrow \sigma$ under environment Γ and b has type σ under environment Γ . If M and N are the terms for a and b respectively, then we use $M N$ for the term of the conclusion.

$$\begin{array}{c}
\text{TERMVAR} \\
\frac{\Gamma \text{ env} \quad (x : \tau) \in \Gamma}{x \Rightarrow x : \Gamma \vdash \tau} \\
\\
\text{TERMLAM} \\
\frac{\Gamma \vdash \tau \text{ type} \quad M \Rightarrow a : \Gamma, (x : \tau) \vdash \sigma}{\lambda(x : \tau) M \Rightarrow \lambda x a : \Gamma \vdash \tau \rightarrow \sigma} \\
\\
\text{TERMAPP} \\
\frac{M \Rightarrow a : \Gamma \vdash \tau \rightarrow \sigma \quad N \Rightarrow b : \Gamma \vdash \tau}{MN \Rightarrow ab : \Gamma \vdash \sigma} \\
\\
\text{TERMPAIR} \\
\frac{M \Rightarrow a : \Gamma \vdash \tau \quad N \Rightarrow b : \Gamma \vdash \sigma}{\langle M, N \rangle \Rightarrow \langle a, b \rangle : \Gamma \vdash \tau \times \sigma} \\
\\
\text{TERMFST} \\
\frac{M \Rightarrow a : \Gamma \vdash \tau \times \sigma}{\text{fst } M \Rightarrow \text{fst } a : \Gamma \vdash \tau} \\
\\
\text{TERMSND} \\
\frac{M \Rightarrow a : \Gamma \vdash \tau \times \sigma}{\text{snd } M \Rightarrow \text{snd } a : \Gamma \vdash \sigma} \\
\\
\text{TERMCOER} \\
\frac{M \Rightarrow a : \Gamma, \Sigma \vdash \tau \quad G \Rightarrow \Gamma \vdash (\Sigma \vdash \tau) \triangleright \sigma}{G \langle M \rangle \Rightarrow a : \Gamma \vdash \sigma}
\end{array}$$

Figure 4.10: System F_t^p term judgment relation

Rule **TERMPAIR** gives type $\tau \times \sigma$ to the pair $\langle a, b \rangle$ under environment Γ if a has type τ under environment Γ and b has type σ under environment Γ . If M and N are the terms for a and b respectively, then we use $\langle M, N \rangle$ for the term of the conclusion. Rule **TERMFST** gives type τ to the first projection $\text{fst } a$ under environment Γ if a has type $\tau \times \sigma$ under environment Γ . If M is the term for the premise, then we write $\text{fst } M$ for the term of the conclusion. Rule **TERMSND** gives type σ to the second projection $\text{snd } a$ under environment Γ if a has type $\tau \times \sigma$ under environment Γ . If M is the term for the premise, then we write $\text{snd } M$ for the term of the conclusion.

Finally, rule **TERMCOER** gives typing $\Gamma \vdash \sigma$ to the term a if it also has typing $\Gamma, \Sigma \vdash \tau$ and there is a coercion from typing $\Sigma \vdash \tau$ to type σ under Γ , written $\Gamma \vdash (\Sigma \vdash \tau) \triangleright \sigma$. If we name M the term witnessing $a : \Gamma, \Sigma \vdash \tau$ and G the coercion proof, then we write $G \langle M \rangle$ the term for the conclusion that a has typing $\Gamma \vdash \sigma$.

We can now define the coercion typing rules which are given on Figure 4.11. The coercion judgment is written $G \Rightarrow \Gamma \vdash (\Sigma \vdash \tau) \triangleright \sigma$. The coercion proof G witnesses the derivation of $\Gamma \vdash (\Sigma \vdash \tau) \triangleright \sigma$. This judgment means that the typing $\Sigma \vdash \tau$ is smaller than the type σ under environment Γ , which can also be seen as the polymorphic type $\forall \Sigma \tau$ being smaller than σ under environment Γ . When the erasable environment Σ is empty, we may write $\Gamma \vdash \tau \triangleright \sigma$.

Rules **COERREFL** and **COERTRANS** are about the closure of the coercion relation by reflexivity and transitivity. The coercion proof \diamond is a witness that type τ is smaller than itself under environment Γ by rule **COERREFL**. Rule **COERTRANS** tells that if coercion G_2 proves that $\Sigma_2 \vdash \tau_2$ is smaller than τ_3 under Γ and coercion G_1 proves that $\Sigma_1 \vdash \tau_1$ is smaller than τ_2 under Γ, Σ_2 then coercion $G_2 \circ G_1$ proves that $\Sigma_2, \Sigma_1 \vdash \tau_1$ is smaller than τ_3 under Γ .

Rule **COERVAR** looks in the environment Γ a coercion hypothesis. In the explicit version this lookup is done using the name given by the coercion variable c . If c is bound to $\tau \triangleright \sigma$ in Γ , then τ is smaller than σ . Rule **COERWEAK** forgets that a coercion G used to extend the environment if the extension is not used by the inner type τ . Concretely, if G witnesses a coercion from τ to σ extending Σ , then $*G$ witnesses a coercion from τ to σ with no environment extension. This is sound since τ is well-formed under Γ by hypothesis.

$$\begin{array}{c}
\text{COERREFL} \\
\diamond \Rightarrow \Gamma \vdash \tau \triangleright \tau
\end{array}
\qquad
\frac{\text{COERTTRANS} \quad G_1 \Rightarrow \Gamma, \Sigma_2 \vdash (\Sigma_1 \vdash \tau_1) \triangleright \tau_2 \quad G_2 \Rightarrow \Gamma \vdash (\Sigma_2 \vdash \tau_2) \triangleright \tau_3}{G_2 \circ G_1 \Rightarrow \Gamma \vdash (\Sigma_2, \Sigma_1 \vdash \tau_1) \triangleright \tau_3}$$

$$\frac{\text{COERVAR} \quad (c : \tau \triangleright \sigma) \in \Gamma}{c \Rightarrow \Gamma \vdash \tau \triangleright \sigma}
\qquad
\frac{\text{COERWEAK} \quad \Gamma, \Sigma \text{ env} \quad G \Rightarrow \Gamma \vdash (\Sigma \vdash \tau) \triangleright \sigma}{*G \Rightarrow \Gamma \vdash \tau \triangleright \sigma}
\qquad
\text{COERTLAM} \quad \Lambda \alpha \Rightarrow \Gamma \vdash (\alpha \vdash \tau) \triangleright \forall \alpha \tau$$

$$\frac{\text{COERTAPP} \quad \Gamma \vdash \sigma \text{ type}}{\cdot \sigma \Rightarrow \Gamma \vdash \forall \alpha \tau \triangleright \tau[\alpha/\sigma]}
\qquad
\frac{\text{COERETAARR} \quad \Gamma \vdash \tau' \text{ type} \quad G_1 \Rightarrow \Gamma, \Sigma \vdash \tau' \triangleright \tau \quad G_2 \Rightarrow \Gamma \vdash (\Sigma \vdash \sigma) \triangleright \sigma'}{G_1 \xrightarrow{\tau'} G_2 \Rightarrow \Gamma \vdash (\Sigma \vdash \tau \rightarrow \sigma) \triangleright \tau' \rightarrow \sigma'}$$

$$\frac{\text{COERETAPROD} \quad G_1 \Rightarrow \Gamma \vdash (\Sigma \vdash \tau) \triangleright \tau' \quad G_2 \Rightarrow \Gamma \vdash (\Sigma \vdash \sigma) \triangleright \sigma'}{G_1 \times G_2 \Rightarrow \Gamma \vdash (\Sigma \vdash \tau \times \sigma) \triangleright \tau' \times \sigma'}
\qquad
\frac{\text{COERBOT} \quad \Gamma \vdash \tau \text{ type}}{\perp \tau \Rightarrow \Gamma \vdash \perp \triangleright \tau}
\qquad
\text{COERTOP} \quad \top \Rightarrow \Gamma \vdash \tau \triangleright \top$$


$$\frac{\text{COERTLAML} \quad \Gamma, \alpha \vdash \tau \text{ type}}{\Lambda(\alpha \triangleleft c : \tau) \Rightarrow \Gamma \vdash (\alpha, (c : \tau \triangleright \alpha) \vdash \rho) \triangleright \forall(\alpha \triangleleft \tau) \rho}
\qquad
\frac{\text{COERTAPPL} \quad \Gamma \vdash \sigma \text{ type} \quad G \Rightarrow \Gamma \vdash \tau[\alpha/\sigma] \triangleright \sigma}{\cdot[\sigma \triangleleft G] \Rightarrow \Gamma \vdash \forall(\alpha \triangleleft \tau) \rho \triangleright \rho[\alpha/\sigma]}$$

$$\frac{\text{COERTLAMU} \quad \Gamma, \alpha \vdash \tau \text{ type}}{\Lambda(\alpha \triangleright c : \tau) \Rightarrow \Gamma \vdash (\alpha, (c : \alpha \triangleright \tau) \vdash \rho) \triangleright \forall(\alpha \triangleright \tau) \rho}
\qquad
\frac{\text{COERTAPPU} \quad \Gamma \vdash \sigma \text{ type} \quad G \Rightarrow \Gamma \vdash \sigma \triangleright \tau[\alpha/\sigma]}{\cdot[\sigma \triangleright G] \Rightarrow \Gamma \vdash \forall(\alpha \triangleright \tau) \rho \triangleright \rho[\alpha/\sigma]}$$

Figure 4.11: System \mathbb{F}_t^p coercion judgment relation

Rules COERTLAM and COERTAPP have to do with polymorphism. The first is type generalization and the second is type instantiation. We define coercion $\Lambda \alpha$ as a proof that typing $\alpha \vdash \tau$ is included in type $\forall \alpha \tau$ under Γ . And we define coercion $\cdot \sigma$ as a proof that we can go from type $\forall \alpha \tau$ to type $\tau[\alpha/\sigma]$ under Γ , given type σ is well-formed under Γ .

Rules COERETAARR and COERETAPROD are about η -expansions so to understand the rules we need to look at their η -expansion views as erasable contexts. The coercion $G_1 \xrightarrow{\tau'} G_2$ can be viewed as the erasable context $\lambda(x : \tau') G_2 \langle \llbracket (G_1(x)) \rrbracket \rangle$. We can display the following typing derivation as a graph:



or as a typing derivation where $\Gamma' = \Gamma, (x : \tau')$:

$$\frac{\frac{\frac{\frac{x \Rightarrow x : \Gamma', \Sigma \vdash \tau' \quad G_1 \Rightarrow \Gamma', \Sigma \vdash \tau' \triangleright \tau}{G_1 \langle x \rangle \Rightarrow x : \Gamma', \Sigma \vdash \tau}}{\boxed{\Rightarrow} : \Gamma', \Sigma \vdash \tau \rightarrow \sigma}}{\boxed{\Rightarrow} (G_1 \langle x \rangle) \Rightarrow \boxed{\Rightarrow} x : \Gamma', \Sigma \vdash \sigma} \quad G_2 \Rightarrow \Gamma' \vdash (\Sigma \vdash \sigma) \triangleright \sigma'}{\frac{G_2 \langle \boxed{\Rightarrow} (G_1 \langle x \rangle) \rangle \Rightarrow \boxed{\Rightarrow} x : \Gamma' \vdash \sigma'}{\lambda(x : \tau') G_2 \langle \boxed{\Rightarrow} (G_1 \langle x \rangle) \rangle \Rightarrow \lambda x \boxed{\Rightarrow} x : \Gamma \vdash \tau' \rightarrow \sigma'}}$$

Rule COERTOP tells that any type is smaller than the top type. And rule COERBOT tells that any well-formed type is bigger than the bottom type.

Finally, rules COERTLAMU, COERTAPPU, COERTLAML, and COERTAPPL have to do with upper and lower bounded polymorphism and are thus similar to rules COERTLAM and COERTAPP. Rule COERTLAMU defines coercion $\Lambda(\alpha \triangleright c : \tau)$ as a proof that typing $\alpha, (c : \alpha \triangleright \tau) \vdash \rho$ is smaller than type $\forall(\alpha \triangleright \tau) \rho$ under Γ . Notice that this coercion binds two variables: the type variable α and the coercion variable c with coercion type $\alpha \triangleright \tau$. Rule COERTAPPU defines the upper bounded instantiation coercion $\cdot[\sigma \triangleright G]$ as a proof that type $\forall(\alpha \triangleright \tau) \rho$ is smaller than $\rho[\alpha/\sigma]$ under environment Γ given σ is well-formed and coercion G is a proof that σ is smaller than $\tau[\alpha/\sigma]$.

Similarly, rule COERTLAML defines coercion $\Lambda(\alpha \triangleleft c : \tau)$ as a proof that typing $\alpha, (c : \tau \triangleright \alpha) \vdash \rho$ is smaller than type $\forall(\alpha \triangleleft \tau) \rho$ under Γ . Notice that this coercion binds two variables: the type variable α and the coercion variable c with coercion type $\tau \triangleright \alpha$. Rule COERTAPPL defines the lower bounded instantiation coercion $\cdot[\sigma \triangleleft G]$ as a proof that type $\forall(\alpha \triangleleft \tau) \rho$ is smaller than $\rho[\alpha/\sigma]$ under environment Γ given σ is well-formed and coercion G is a proof that $\tau[\alpha/\sigma]$ is smaller than σ .

It remains to define the well-formedness rules. A type τ is well-formed under environment Γ , written $\Gamma \vdash \tau$ **type**, if it has a derivation using the rules given in Figure 4.12. Similarly, an environment Γ is well-formed, written Γ **env**, if it has a derivation using the rules given in the same figure.

Rule TYPEVAR tells that a type variable α is well-formed under environment Γ if it is bound in the environment. Since we want to extract from a type derivation that the environment is well-formed, we need to ask the environment Γ to be well-formed. Rule TYPEARR tells that type $\tau \rightarrow \sigma$ is well-formed under environment Γ , if both τ and σ are well-formed under Γ . Similarly, type $\tau \times \sigma$ is well-formed under environment Γ by rule TYPEPROD, if both τ and σ are well-formed under Γ .

Rule TYPEFOR tells that the polymorphic type $\forall \alpha \tau$ is well-formed under environment Γ if type τ is well-formed under the extended environment Γ, α . The top and bottom types are well-formed, according to rules TYPETOP and TYPEBOT respectively, if their environment is well-formed. Rules TYPEFORU and TYPEFORL are similar. Type $\forall(\alpha \triangleright \tau) \rho$ is well-formed under environment Γ if type ρ is well-formed under the extended environment $\Gamma, \alpha, (c : \alpha \triangleright \tau)$. Similarly, type $\forall(\alpha \triangleleft \tau) \rho$ is well-formed under environment Γ if type ρ is well-formed under the extended environment $\Gamma, \alpha, (c : \tau \triangleright \alpha)$.

Since bindings are related to abstractions and not variables, and since System F_t^p has abstractions binding more than one variable, the well-formedness derivation of an environment has to split the environment as a list of abstractions. System F_t^p has four kind of abstractions: term abstractions with binding $(x : \tau)$, type abstraction with binding α , upper bounded type abstraction with binding $\alpha, (c : \alpha \triangleright \tau)$, and lower bounded type abstraction with binding $\alpha, (c : \tau \triangleright \alpha)$.

$$\begin{array}{c}
\text{TYPEVAR} \\
\frac{\Gamma \text{ env} \quad \alpha \in \Gamma}{\Gamma \vdash \alpha \text{ type}} \\
\\
\text{TYPEARR} \\
\frac{\Gamma \vdash \tau \text{ type} \quad \Gamma \vdash \sigma \text{ type}}{\Gamma \vdash \tau \rightarrow \sigma \text{ type}} \\
\\
\text{TYPEPROD} \\
\frac{\Gamma \vdash \tau \text{ type} \quad \Gamma \vdash \sigma \text{ type}}{\Gamma \vdash \tau \times \sigma \text{ type}} \\
\\
\text{TYPEFOR} \quad \text{TYPEBOT} \quad \text{TOTYPETOP} \quad \text{TYPEFORL} \\
\frac{\Gamma, \alpha \vdash \tau \text{ type}}{\Gamma \vdash \forall \alpha \tau \text{ type}} \quad \frac{\Gamma \text{ env}}{\Gamma \vdash \perp \text{ type}} \quad \frac{\Gamma \text{ env}}{\Gamma \vdash \top \text{ type}} \quad \frac{\Gamma, \alpha \vdash \tau \text{ type} \quad \Gamma, \alpha, (c : \tau \triangleright \alpha) \vdash \rho \text{ type}}{\Gamma \vdash \forall (\alpha \triangleleft \tau) \rho \text{ type}} \\
\\
\text{TYPEFORU} \\
\frac{\Gamma, \alpha \vdash \tau \text{ type} \quad \Gamma, \alpha, (c : \alpha \triangleright \tau) \vdash \rho \text{ type}}{\Gamma \vdash \forall (\alpha \triangleright \tau) \rho \text{ type}} \\
\\
\text{ENVEMPTY} \quad \text{ENVTERM} \quad \text{ENVTYPE} \\
\frac{}{\emptyset \text{ env}} \quad \frac{x \notin \text{dom}(\Gamma) \quad \Gamma \vdash \tau \text{ type}}{\Gamma, (x : \tau) \text{ env}} \quad \frac{\Gamma \text{ env} \quad \alpha \notin \text{dom}(\Gamma)}{\Gamma, \alpha \text{ env}} \\
\\
\text{ENVTYPEL} \quad \text{ENVTYPEU} \\
\frac{c, \alpha \notin \text{dom}(\Gamma) \quad \Gamma, \alpha \vdash \tau \text{ type}}{\Gamma, \alpha, (c : \tau \triangleright \alpha) \text{ env}} \quad \frac{c, \alpha \notin \text{dom}(\Gamma) \quad \Gamma, \alpha \vdash \tau \text{ type}}{\Gamma, \alpha, (c : \alpha \triangleright \tau) \text{ env}}
\end{array}$$

Figure 4.12: System F_t^p well-formedness relations

Once the environment is split by abstractions we proceed by recurrence over this list of abstractions. Either the list is empty and we have the empty environment \emptyset which is always well-formed by rule `ENVEMPTY`. Or we have a non-empty list of abstractions and we look at the last one. The remaining list has to be well-formed and the last abstraction has to be well-formed under the remaining environment. A term abstraction associating the term variable x to the type τ is well-formed by rule `ENVTERM` if the type τ is well-formed under the remaining environment Γ and the term variable x is not already bound in Γ . A type abstraction α is well-formed by rule `ENVTYPE` if the type variable α is not already bound in the remaining environment. Upper and lower type abstractions, $\alpha, (c : \alpha \triangleright \tau)$ and $\alpha, (c : \tau \triangleright \alpha)$ respectively, are well-formed by rules `ENVTYPEU` and `ENVTYPEL` respectively if the type variable α and coercion variable c are not already bound in the remaining environment Γ and if the type τ is well-formed under the environment Γ, α since we allow recursive bounds.

4.4 Properties

We first describe the properties linking the implicit version and the explicit version of the type system as we did in Chapter 3. We then prove the strong normalization of the explicit reduction of well-typed terms by translating our judgments in System F . Using termination we prove the confluence of the explicit reduction. We then show how the explicit reduction corresponds to the implicit reduction in the bisimulation lemma. And we finally prove that System F_t^p is sound and strongly normalizing in both its explicit and implicit versions. We can conclude from these properties that well-typed terms in System F_t^p strongly normalize to a unique value without encountering any error.

$$\begin{array}{lll}
[x] = x & \llbracket \langle M, N \rangle \rrbracket = \langle \llbracket M \rrbracket, \llbracket N \rrbracket \rangle & \\
\llbracket \lambda(x : \tau) M \rrbracket = \lambda x \llbracket M \rrbracket & \llbracket \text{fst } M \rrbracket = \text{fst } \llbracket M \rrbracket & \llbracket G \langle M \rangle \rrbracket = \llbracket M \rrbracket \\
\llbracket M N \rrbracket = \llbracket M \rrbracket \llbracket N \rrbracket & \llbracket \text{snd } M \rrbracket = \text{snd } \llbracket M \rrbracket &
\end{array}$$

Figure 4.13: System F_t^p drop function

4.4.1 Implicit vs. Explicit version

We first give the expected properties of an explicit type system. A judgment has to be unique according to its explicit entity. For instance, when $M \Rightarrow a : \Gamma \vdash \tau$ holds, then a is determined by M , and τ is a function of the term M and the environment Γ . For coercions, when we have $G \Rightarrow \Gamma \vdash (\Sigma \vdash \tau) \triangleright \sigma$, then the erasable environment Σ is a function of coercion G and environment Γ , and the type σ is a function of coercion G , environment Γ , and type τ .

Lemma 43 (Uniqueness). *The following assertions hold.*

- If $M \Rightarrow a_1 : \Gamma_1 \vdash \tau_1$ and $M \Rightarrow a_2 : \Gamma_2 \vdash \tau_2$ hold, then $a_1 = a_2$ holds.
- If $M \Rightarrow a : \Gamma \vdash \tau_1$ and $M \Rightarrow a : \Gamma \vdash \tau_2$ hold, then $\tau_1 = \tau_2$ hold.
- If $G \Rightarrow \Gamma \vdash (\Sigma_1 \vdash \tau_1) \triangleright \sigma_1$ and $G \Rightarrow \Gamma \vdash (\Sigma_2 \vdash \tau_2) \triangleright \sigma_2$ hold, then $\Sigma_1 = \Sigma_2$ holds.
- If $G \Rightarrow \Gamma \vdash (\Sigma \vdash \tau) \triangleright \sigma_1$ and $G \Rightarrow \Gamma \vdash (\Sigma \vdash \tau) \triangleright \sigma_2$ hold, then $\sigma_1 = \sigma_2$ holds.

Proof. For each assertion, by induction on the first hypothesis and inversion of the second. The inversion leads to exactly one rule, which is actually the same as the one for the induction. Rule `TERMVAR` uses the fact that a well-formed environment binds each term variable at most once. All other cases simply use induction hypotheses. \square

Actually a is a function of M even if M is not well-typed. This is not useful as a result, but the function in question is useful for the bisimulation properties. We call this function the *drop* function and we write it $\llbracket M \rrbracket$. It is simply defined by dropping the annotations in M . The formal definition is given on Figure 4.13. This lemma explains why we usually omit a in explicit term judgments and write $M : \Gamma \vdash \tau$ instead of $M \Rightarrow a : \Gamma \vdash \tau$.

Lemma 44. *If $M \Rightarrow a : \Gamma \vdash \tau$ holds, then $a = \llbracket M \rrbracket$ holds.*

Proof. By induction. \square

The next lemma tells that M is actually contained in the implicit typing derivation. In other words, from a derivation of $a : \Gamma \vdash \tau$ we can extract the term M such that $M \Rightarrow a : \Gamma \vdash \tau$ holds. The environment Γ has to be filled with distinct coercion variables. Reciprocally, if $M \Rightarrow a : \Gamma \vdash \tau$ holds, we can extract a derivation of $a : \Gamma' \vdash \tau$ where Γ' is obtained from Γ by removing coercion variables. Similarly for coercions, we can add or remove the witness.

Lemma 45 (Equivalence). *The following assertions hold.*

- $a : \Gamma \vdash \tau$ holds if and once if $M \Rightarrow a : \Gamma \vdash \tau$ holds for some M .
- $\Gamma \vdash (\Sigma \vdash \tau) \triangleright \sigma$ holds if and only if $G \Rightarrow \Gamma \vdash (\Sigma \vdash \tau) \triangleright \sigma$ holds for some G .

Proof. By induction. The reciprocals are very easy since we only remove information. The implicit to explicit direction simply uses the induction hypotheses and the syntax of the side judgments. \square

$$\begin{array}{lcl}
\widehat{x} = x & \langle \widehat{M}, \widehat{N} \rangle = \langle \widehat{M}, \widehat{N} \rangle & \widehat{G\langle M \rangle} = \widehat{G}[\widehat{M}] \\
\lambda(x : \tau) \widehat{M} = \lambda(x : \widehat{\tau}) \widehat{M} & \widehat{\text{fst } M} = \text{fst } \widehat{M} & \\
\widehat{M N} = \widehat{M} \widehat{N} & \widehat{\text{snd } M} = \text{snd } \widehat{M} &
\end{array}$$

Figure 4.14: System F_t^p term reification function

$$\begin{array}{lcl}
\widehat{\alpha} = \alpha & \widehat{\top} = \forall \alpha \alpha \rightarrow \alpha & \\
\widehat{\tau \rightarrow \tau} = \widehat{\tau} \rightarrow \widehat{\sigma} & \widehat{\perp} = \forall \alpha \alpha & \\
\widehat{\tau \times \tau} = \widehat{\tau} \times \widehat{\sigma} & \forall (\widehat{\alpha \triangleright \tau}) \rho = \forall \alpha (\alpha \rightarrow \widehat{\tau}) \rightarrow \widehat{\rho} & \\
\widehat{\forall \alpha \tau} = \forall \alpha \widehat{\tau} & \forall (\widehat{\alpha \triangleleft \tau}) \rho = \forall \alpha (\widehat{\tau} \rightarrow \alpha) \rightarrow \widehat{\rho} &
\end{array}$$

Figure 4.15: System F_t^p type reification function

4.4.2 Termination

Termination for the explicit version is shown by reification into System F . We reify types, environments, terms, and coercion proofs into System F . We show that the reification of System F_t^p derivations are also valid derivations in System F . So the reification of a well-typed explicit term of System F_t^p remains well-typed in System F . We also show a simulation from the explicit reduction in System F_t^p to the reduction of System F . Thus a non-terminating well-typed explicit term in System F_t^p implies the existence of a non-terminating well-typed term in System F , which is known to be strongly normalizing. We use a hat to denote the reification of an object.

We give the reification of explicit terms in Figure 4.14 and write \widehat{M} for the reification of the term M . Terms of System F_t^p are reified to terms of System F . Variables are reified to variables, abstractions to abstractions, applications to applications, pairs to pairs, and projections to projections, and recursively for subterms. The only interesting case is the coercion application $G\langle M \rangle$. Coercions of System F_t^p are reified to System F multi-hole (but with at least one hole, see the reification of coercions) contexts. So the reification of the coercion application of G to the term M is the application of the context \widehat{G} to the term \widehat{M} , leading to a possible duplication of \widehat{M} .

We give the reification of types in Figure 4.15. Type variables, arrow types, product types, and polymorphic types are simply translated to their equivalent in System F . The top type is reified to the polymorphic identity type. The bottom type is reified to the polymorphic bottom type $\forall \alpha \alpha$. Upper and lower bounded polymorphic types are reified to polymorphic arrow types: abstract coercions are reified to functions. Upper bounded polymorphism reification first abstracts over the type variable α , then over the function $\alpha \rightarrow \widehat{\tau}$, with body $\widehat{\rho}$. Lower bounded polymorphism is similar, but since τ is a lower bound, the reified abstract coercion is a function from $\widehat{\tau}$ to α .

The most interesting reification function is for coercions given in Figure 4.16. Coercions are reified to multi-hole contexts, with at least one hole. We write id for the identity function $\lambda(x : \tau) x$ of System F at the correct type τ depending of where it is used. There is always exactly one type, so it lightens the notations. We use id to add extra reduction steps for the simulation to hold. We also use top for $\lambda(y : \tau) \Lambda \alpha \lambda(x : \alpha) x$ where τ is the type where top

$$\begin{array}{l}
\widehat{c} = x_c \ [] \\
\widehat{\diamond} = \text{id} \ [] \\
\widehat{G_2 \circ G_1} = \text{id} (\widehat{G_2}[\widehat{G_1}]) \\
\widehat{*G} = \text{id} \widehat{G} \\
\widehat{\Lambda\alpha} = \Lambda\alpha \ [] \\
\widehat{\cdot\sigma} = \llbracket [\widehat{\sigma}] \rrbracket \\
\widehat{G_1 \xrightarrow{\tau'} G_2} = \lambda(x : \widehat{\tau}') \widehat{G_2}[\llbracket \widehat{G_1}[x] \rrbracket] \\
\widehat{G_1 \times G_2} = \langle \widehat{G_1}[\text{fst} \ []], \widehat{G_2}[\text{snd} \ []] \rangle
\end{array}
\qquad
\begin{array}{l}
\widehat{\top} = \text{top} \ [] \\
\widehat{\perp\tau} = \llbracket [\widehat{\tau}] \rrbracket \\
\Lambda(\alpha \triangleright c : \tau) = \Lambda\alpha \lambda(x_c : \alpha \rightarrow \widehat{\tau}) \ [] \\
\cdot[\widehat{\sigma} \triangleright G] = \llbracket [\widehat{\sigma}] (\lambda(x : \widehat{\sigma}) \widehat{G}[x]) \rrbracket \\
\Lambda(\alpha \triangleleft c : \tau) = \Lambda\alpha \lambda(x_c : \widehat{\tau} \rightarrow \alpha) \ [] \\
\cdot[\widehat{\sigma} \triangleleft G] = \llbracket [\widehat{\sigma}] (\lambda(x : \widehat{\tau}[\alpha/\widehat{\sigma}]) \widehat{G}[x]) \rrbracket
\end{array}$$

Figure 4.16: System F_t^p coercion reification function

$$\begin{array}{l}
\widehat{\emptyset} = \emptyset \\
\widehat{\Gamma, \alpha} = \widehat{\Gamma}, \alpha \\
\widehat{\Gamma, (x : \tau)} = \widehat{\Gamma}, (x : \widehat{\tau}) \quad \widehat{\Gamma, (c : \tau \triangleright \sigma)} = \widehat{\Gamma}, (x_c : \widehat{\tau} \rightarrow \widehat{\sigma})
\end{array}$$

Figure 4.17: System F_t^p environment reification function

is used, similarly to **id**. **top** is used to forget the type of an expression, but still keep it under hand to reduce inside it for the simulation.

Coercion variables are reified to term variables applied to the hole. We actually partition the term variables of System **F** into two parts: one part for the reification of System F_t^p term variables, and one for coercion variables. The reflexivity coercion applies the identity function to the hole in order to have an extra reduction step. The transitivity coercion composes the two sub-contexts and applies the identity function to add one reduction step. Type abstraction and application are reified to their equivalent in System **F**. Arrow and product η -expansions are reified as computational η -expansions. Arrow η -expansion is $\lambda x \ [] x$ where we add the two sub-contexts for the argument and body. Product η -expansion is $\langle \text{fst} \ [], \text{snd} \ [] \rangle$. Similarly we add the sub-contexts for the first and second components. Notice that this is the only coercion where we create more than one hole. This is not a problem since we only need a forward simulation from System F_t^p to System **F**.

The top coercion is reified to the hole applied to **top**, which is a way to keep the hole under hand while the final type is the polymorphic identity type, which we use as the reification for **top**. The bottom instantiation coercion is reified to System **F** type instantiation. Upper bounded type abstraction $\Lambda(\alpha \triangleright c : \tau)$ is reified to a type abstraction $\Lambda\alpha$ followed by a function abstraction $\lambda(x_c : \alpha \rightarrow \widehat{\tau})$ for the abstract coercion $(c : \alpha \triangleright \tau)$. Conversely, upper bounded type application $\cdot[\sigma \triangleright G]$ is a type application $[\widehat{\sigma}]$ followed by the application of the function reifying the coercion $\lambda(x : \widehat{\sigma}) \widehat{G}[x]$. To build this function, we close the context with a term abstraction and use the variable to fill the holes of the context. The type annotation of the term abstraction is the type of the hole of \widehat{G} which is unique. Lower bounded abstraction and application are similar.

Finally, environments are reified in Figure 4.17. The empty environment is reified to the empty environment. Term bindings are reified to term bindings, type bindings to type bindings, and coercion bindings $(c : \tau \triangleright \sigma)$ to term bindings $(x_c : \widehat{\tau} \rightarrow \widehat{\sigma})$ because coercions are reified to functions.

We write $(J)_s$ for judgments in System F_t^p , because it is the source language. And we

write $(J)_t$ for the judgments of System F, because it is the target judgment. We show that derivations of System F_t^p are reified to valid derivations of System F.

Lemma 46 (Typing preservation). *The following properties hold.*

- If $(\Gamma \text{ env})_s$ holds, then $(\widehat{\Gamma} \text{ env})_t$ holds.
- If $(\Gamma \vdash \tau \text{ type})_s$ holds, then $(\widehat{\Gamma} \vdash \widehat{\tau} \text{ type})_t$ holds.
- If $(M \Rightarrow a : \Gamma \vdash \tau)_s$ holds, then $(\widehat{M} \Rightarrow \widehat{\Gamma} \vdash a : \widehat{\tau})_t$ holds.
- If $(G \Rightarrow \Gamma \vdash (\Sigma \vdash \tau) \triangleright \sigma)_s$ and $(M \Rightarrow \widehat{\Gamma}, \widehat{\Sigma} \vdash a : \widehat{\tau})_t$ hold, then $(\widehat{G}[M] \Rightarrow \widehat{\Gamma} \vdash a : \widehat{\sigma})_t$ holds.

Proof. By induction. The only interesting assertion is the last one. Reflexivity, transitivity, and coercion variable are simple context transformations. The weakening coercion uses the weakening lemma of System F. Type abstraction and application use their homologous rule.

For rule COERETAARR, we first abstract over a fresh term variable x . We then use the induction hypothesis for the body context, which we feed with the application of the hypothesis weakened with the term binding of x , to x under the context of the argument coercion. Rule COERETAPROD is similar without the weakening part because there is no binding.

Bottom instantiation uses the type instantiation rule. The top coercion uses the typing derivation of **top**. Upper and lower bounded polymorphic rules are simple uses of type and term abstractions and applications. \square

Lemma 47 (Simulation). *If $M \rightsquigarrow_{\beta_L} N$ holds, then $\widehat{M} \rightsquigarrow^+ \widehat{N}$ holds.*

Proof. By induction. Rule REDCTX is obvious for all contexts but coercion application. We know that coercions are reified to multi-hole contexts with at least one hole. So we need to repeat the steps of the induction hypothesis as many times as there are holes. Rules REDAPP, REDFST, and REDSND are reified to their analog rules.

Rule REDREFL is reified to $\text{id } \widehat{M}$ which reduces to \widehat{M} . Without **id**, it would not have been possible to do a reduction step in System F. Rule REDTRANS and REDWEAK use the same mechanism. Rule REDPOLY reifies to the analog reduction rule.

For rules REDETAARR and REDETAAPROD we reduce the inner redexes, where the hole was. Because both rules reduce under coercion contexts, they may duplicate reduction steps. Rules REDPOLYL and REDPOLYU use type and term abstraction reduction rules. \square

Lemma 48 (Termination). *If $M \Rightarrow a : \Gamma \vdash \tau$ holds, then M strongly normalizes.*

Proof. \widehat{M} is well-typed in System F by Lemma 46. Using System F strong normalization result, we know that reduction of \widehat{M} terminates. If we had an infinite reduction path for M , then by Lemma 47 we would have one for \widehat{M} too, which is a contradiction. \square

4.4.3 Confluence

The explicit reduction is locally confluent and thus confluent since reduction terminates.

Lemma 49 (Local confluence). *If $M \rightsquigarrow_{\beta_L} M_1$ and $M \rightsquigarrow_{\beta_L} M_2$ hold, then there is a N such that $M_1 \rightsquigarrow_{\beta_L}^* N$ and $M_2 \rightsquigarrow_{\beta_L}^* N$ hold.*

Proof. There are no critical pairs. \square

Corollary 50 (Confluence). *If $M \rightsquigarrow_{\beta\iota}^* M_1$ and $M \rightsquigarrow_{\beta\iota}^* M_2$ hold, then there is a N such that $M_1 \rightsquigarrow_{\beta\iota}^* N$ and $M_2 \rightsquigarrow_{\beta\iota}^* N$ hold.*

Proof. By Newman's lemma and Lemma 49 and 48. \square

4.4.4 Bisimulation

This section makes the link between the λ -calculus reduction and the explicit reduction. The main result is the bisimulation lemma. It helps to show the soundness of the implicit version but has also its own meaning. In a few words, the bisimulation lemma guarantees that typing annotations of the explicit term do not alter or block the reduction of its implicit underlying term. In other words, typing annotations are erasable (see Section 6.1.5 for non-erasability).

In order to show the bisimulation lemma, in particular the backward simulation, we need to show a classification lemma on ι -normal forms. This lemma also shows consistency of coercions. Only the first assertion of the iota classification lemma is used for bisimulation, and only the second assertion is used for consistency. The remaining assertions are used for the induction. We define retyping contexts Q as sequences of coercion applications: $Q ::= [] \mid G\langle Q \rangle$.

Lemma 51 (Iota classification). *If $Q[\lambda(x : \rho) M] \Rightarrow \lambda x a : \Gamma \vdash \tau$ (resp. $Q[\langle M, N \rangle] \Rightarrow \langle a, b \rangle : \Gamma \vdash \tau$) holds and $Q[\lambda(x : \rho) M]$ (resp. $Q[\langle M, N \rangle]$) is in ι -normal form, then the following assertions hold.*

- *If τ is $\tau_1 \rightarrow \tau_2$ (resp. $\tau_1 \times \tau_2$), then $Q = []$.*
- *τ is not $\tau_1 \times \tau_2$ (resp. $\tau_1 \rightarrow \tau_2$).*
- *If τ is $\forall \alpha \tau'$, then $Q = \Lambda \alpha \langle Q' \rangle$.*
- *If τ is $\forall (\alpha \triangleright \sigma) \tau'$, then $Q = \Lambda (\alpha \triangleright c : \sigma) \langle Q' \rangle$.*
- *If τ is $\forall (\alpha \triangleleft \sigma) \tau'$, then $Q = \Lambda (\alpha \triangleleft c : \sigma) \langle Q' \rangle$.*
- *τ is not \perp .*
- *For all $\alpha, (c : \alpha \triangleright \sigma)$ in Γ , τ is not α .*

Proof. We detail the arrow case. The product case is similar. We proceed by induction on Q . If Q is a hole, then all assertions hold. Let's proceed by cases on G when Q is of the form $G\langle Q \rangle$.

Reflexivity, transitivity, and weakening contradict the ι -normal form hypothesis. Type abstraction satisfies all assertions. If $Q = \cdot \sigma \langle Q' \rangle$, then $Q'[\lambda(x : \rho) M]$ is in ι -normal form and has type $\forall \alpha \tau'$ for some type τ' . By induction hypothesis $Q' = \Lambda \alpha \langle Q'' \rangle$ which contradicts the ι -normal form hypothesis.

For the arrow η -expansion, we have $Q = (G_1 \xrightarrow{\tau'} G_2) \langle Q' \rangle$, so by typing and induction hypothesis we have $Q' = []$. This contradicts the ι -normal form hypothesis. Similarly for product η -expansion. The bottom instantiation is impossible by typing and induction hypothesis. The

top coercion satisfies all assertions. Upper and lower bounded abstractions satisfy all assertions, while upper and lower bounded applications break the ι -normal form hypothesis. These cases are similar to polymorphism.

Finally, coercion variables rely on the fact that the environment Γ is well-formed. For upper bounded coercion variables ($c : \alpha \triangleright \sigma$), there is a contradiction using the induction hypothesis, because $Q = c\langle Q' \rangle$ and Q' cannot end with type α . For lower bounded coercion variables ($c : \sigma \triangleright \alpha$), all assertions are satisfied. The only interesting assertion is the last one. By the well-formedness of Γ , we know that α is bound just before ($c : \sigma \triangleright \alpha$). Because α cannot be bound more than once, there is no $\alpha, (c' : \alpha \triangleright \sigma')$ bound in Γ . \square

Notice that the lemma does not require the whole term $Q[\lambda(x : \rho) M]$ to be ι -irreducible. The proof only needs the retyping context Q and its interaction with the term abstraction, $Q[\lambda(x : \rho) \cdot]$, to be in ι -normal form. The term M may contain ι -redexes, the proof does not inspect M .

The forward simulation means that β -reduction corresponds to the λ -calculus reduction and that ι -reduction is only static: it does not change the computational behavior of the term. In other words β -steps are steps of the λ -calculus and ι -steps are erasable steps.

Lemma 52 (Forward simulation). *The following assertions hold.*

- *If $M \rightsquigarrow_{\beta} N$ holds, then $\lfloor M \rfloor \rightsquigarrow \lfloor N \rfloor$ holds.*
- *If $M \rightsquigarrow_{\iota} N$ holds, then $\lfloor M \rfloor = \lfloor N \rfloor$ holds.*
- *If the well-typed term M is an explicit value, then $\lfloor M \rfloor$ is an implicit value.*

Proof. By induction for the first two assertions. For the last one, it suffices to see that prevalues drop on prevalues and values drop on values with one exception: values coerced with a coercion variable are prevalues. If $\lfloor c\langle v \rangle \rfloor$ is used as an implicit prevalue and $\lfloor v \rfloor$ is not a prevalue (it is either an abstraction or a pair), then by Lemma 51 we have a contradiction. The explicit value is actually of the form $E[Q_1[c\langle Q_2[M] \rangle]]$, where M is a computational constructor as $\lfloor v \rfloor$ is not a prevalue, and E ends with a computational destructor because $\lfloor c\langle v \rangle \rfloor$ is used as a prevalue. We use Lemma 51 with $Q_1[c\langle Q_2[M] \rangle]$. If the computational destructor and constructor correspond, then we are in the first assertion and $Q_1[c\langle Q_2[\] \rangle]$ is empty. This is impossible because there is at least c . If the computational destructor and constructor do not correspond, then we are in the second assertion. Again, this is impossible. \square

Reciprocally, the backward simulation tells that not only β -reduction steps are λ -calculus steps, but that they contain all possible λ -calculus steps modulo some erasable ι -reduction steps. This means that coercions cannot block a λ -calculus redex. For instance, the following term $((\cdot \text{Int} \circ \Lambda \alpha) \langle \lambda(x : \alpha) x \rangle) 49$, which is the polymorphic identity instantiated with `Int` and applied to `49`, drops to $(\lambda x x) 49$ which reduces to `49`. However, in order to have access to this redex in the source term, we need to do two ι -steps: `REDTRANS` and `REDPOLY`.

Lemma 53 (Backward simulation). *If $\lfloor M \rfloor \rightsquigarrow b$ holds and M is well-typed, then there is an N such that $M \rightsquigarrow_{\iota}^* \rightsquigarrow_{\beta} N$ and $b = \lfloor N \rfloor$ hold.*

Proof. By Lemma 48 we can assume M in ι -normal form. We proceed by induction. If the reduction rule is `REDCTX`, we know that M is of the form $G\langle E[M'] \rangle$ with $\lfloor M' \rfloor \rightsquigarrow b'$. By

induction hypothesis we find N' and use $G\langle E[N'] \rangle$. We do rule REDAPP in details. Rules REDFST and REDSND are similar.

We know that $\lfloor M \rfloor = (\lambda x a_1) a_2$ which reduces on $a_1[x/a_2]$. By inversion of the drop function, we know that M is of the form $E[Q[\lambda(x : \tau_2) M_1] M_2]$ with $\lfloor M_1 \rfloor = a_1$ and $\lfloor M_2 \rfloor = a_2$. By typing we know that $Q[\lambda(x : \tau_2) M_1]$ has type $\tau_2 \rightarrow \tau_1$. By Lemma 51 we know that Q is $\llbracket \cdot \rrbracket$. Hence M contains a redex the term abstraction is right under the application and reduces to $E[M_1[x/M_2]]$ which drops on $a_1[x/a_2]$ which concludes. \square

4.4.5 Soundness

We prove soundness of the explicit and implicit type systems with the usual subject reduction and progress lemmas. However, the subject reduction and progress lemmas of the implicit version are proved using the subject reduction and progress lemmas of the explicit version and the bisimulation lemma. In order to do these syntactical proofs, we need to prove the usual syntactical lemma about weakening and substitution.

The weakening lemma tells that if an object is well-formed (resp. well-typed) under environment Γ , then it is also well-formed (resp. well-typed) under an extended well-formed environment Γ' .

Lemma 54 (Weakening). *If $\Gamma \subseteq \Gamma'$ and Γ' env hold, then the following assertions hold:*

- *If $\Gamma \vdash \tau$ type holds, then $\Gamma' \vdash \tau$ type holds.*
- *If $G \Rightarrow \Gamma \vdash (\Sigma \vdash \tau) \triangleright \sigma$ holds and $\text{dom}(\Sigma)$ is disjoint from $\text{dom}(\Gamma')$, then Γ', Σ env and $G \Rightarrow \Gamma' \vdash (\Sigma \vdash \tau) \triangleright \sigma$ hold.*
- *If $M \Rightarrow a : \Gamma \vdash \tau$ holds, then $M \Rightarrow a : \Gamma' \vdash \tau$ holds.*

Proof. By mutual induction. We only give the non-trivial and non-similar cases. For rule TYPEVAR, if $\alpha \in \Gamma$ then $\alpha \in \Gamma'$ too. For rule TYPEFORL, because α could have been renamed we have that environment Γ', α is well-formed and by induction hypothesis that $\Gamma', \alpha \vdash \tau$ type holds. We deduce that $\Gamma', \alpha, (c : \tau \triangleright \alpha)$ holds and by induction hypothesis that $\Gamma', \alpha, (c : \tau \triangleright \alpha) \vdash \rho$ type holds. For rule TERMCOR, we may rename the variables bound in Σ in order to avoid the domain of Γ' . We first use the induction hypothesis on the coercion since we need Γ', Σ env to use the induction hypothesis on the term judgment. \square

If a type σ is well formed under environment Γ and τ is well-formed under $\Gamma, \alpha, C, \Gamma'$ (where C is a potential coercion binding of α), then $\tau[\alpha/\sigma]$ is well-formed under $\Gamma, \Gamma'[\alpha/\sigma]$. This lemma is called type substitution. The coercion binding C is maximally inserted, which means Γ' cannot start with a coercion binding. Notice that we don't ask for $\Gamma \vdash C[\alpha/\sigma]$ to hold because coercion variables are not used in type well-formedness judgment.

Lemma 55 (Type substitution). *If $\Gamma \vdash \sigma$ type holds, then the following assertions hold.*

- *If $\Gamma, \alpha, C, \Gamma'$ env holds, then $\Gamma, \Gamma'[\alpha/\sigma]$ env holds.*
- *If $\Gamma, \alpha, C, \Gamma' \vdash \tau$ type holds, then $\Gamma, \Gamma'[\alpha/\sigma] \vdash \tau[\alpha/\sigma]$ type holds.*

Proof. By induction. The only interesting rule is TYPEVAR. There are two cases depending on whether the type variable to instantiate is the same as the one of the judgment. If we have $\Gamma, \alpha, \Gamma' \vdash \alpha$ type then we have to show that $\Gamma, \Gamma'[\alpha/\sigma] \vdash \sigma$ type holds which we do by

Lemma 54 and induction hypothesis. If we have $\Gamma, \alpha, \Gamma' \vdash \beta$ **type** for $\beta \neq \alpha$, then by induction hypothesis we have $\Gamma, \Gamma'[\alpha/\sigma] \vdash \beta$ **type**. \square

The extraction lemma tells when the sub-judgments of a well-formed judgment are also well-formed. From a type derivation we can extract that the environment is well-formed. From a coercion derivation, using the hypothesis that its left type is well-formed, we can extract the well-formedness of its right type. Finally, from a term derivation we can extract that its type is well-formed under its environment, and hence that the environment itself is well-formed.

Lemma 56 (Extraction). *The following assertions hold.*

- If $\Gamma \vdash \tau$ **type holds**, then Γ **env holds**.
- If $\Gamma \vdash (\Sigma \vdash \tau) \triangleright \sigma$ and $\Gamma, \Sigma \vdash \tau$ **type hold**, then $\Gamma \vdash \sigma$ **type holds**.
- If $a : \Gamma \vdash \tau$ **holds**, then $\Gamma \vdash \tau$ **type holds**.

Proof. By mutual induction. We detail the cases where the conclusion is not in hypothesis (like rule **TYPEVAR**), nor obtained by inversion of the induction hypothesis (like rule **TYPEFOR**). For rule **COERVAR**, we have that Γ is well-formed by hypothesis from which we extract by Lemma 54 that σ is well-formed under it. For rule **COERWEAK** we use Lemma 54 to call the induction hypothesis. For rules **COERTAPP**, **COERTAPPL**, and **COERTAPPU**, we use Lemma 55. \square

Term substitution tells that the substitution $a[x/b]$ has type σ under environment Γ if the argument b has type τ under environment Γ and the body a has type σ under environment Γ extended with x associated to τ .

Lemma 57 (Term substitution). *If $N \Rightarrow b : \Gamma \vdash \tau$ hold, then the following assertions hold.*

- If $\Gamma, (x : \tau), \Gamma'$ **env holds**, then Γ, Γ' **env holds**.
- If $\Gamma, (x : \tau), \Gamma' \vdash \sigma$ **type holds**, then $\Gamma, \Gamma' \vdash \sigma$ **type holds**.
- If $G \Rightarrow \Gamma, (x : \tau), \Gamma' \vdash (\Sigma \vdash \sigma) \triangleright \sigma'$ **holds**, then $G \Rightarrow \Gamma, \Gamma' \vdash (\Sigma \vdash \sigma) \triangleright \sigma'$ **holds**.
- If $M \Rightarrow a : \Gamma, (x : \tau), \Gamma' \vdash \sigma$ **holds**, then $M[x/N] \Rightarrow a[x/b] : \Gamma, \Gamma' \vdash \sigma$ **holds**.

Proof. By mutual induction. For rule **TERMVAR**, we use Lemma 54 when the term variables correspond. All other rules use induction hypotheses. \square

If a type σ and a coercion G' are well formed under environment Γ , then for all well-formed judgments under $\Gamma, \alpha, (c : \tau_1 \triangleright \tau_2), \Gamma'$ (where $(c : \tau_1 \triangleright \tau_2)$ is a potential coercion binding), the same judgment after the type substitution $[\alpha/\sigma]$ and coercion substitution $[c/G']$ is well-formed under $\Gamma, \Gamma'[\alpha/\sigma]$. This lemma is called bounded type substitution. The coercion binding $(c : \tau_1 \triangleright \tau_2)$ is maximally inserted, which means that Γ' cannot start with a coercion binding.

Lemma 58 (Bounded type substitution). *If $\Gamma \vdash \sigma$ **type** and $G' \Rightarrow \Gamma \vdash \tau_1[\alpha/\sigma] \triangleright \tau_2[\alpha/\sigma]$ **hold**, then the following assertions hold.*

- If $G \Rightarrow \Gamma, \alpha, (c : \tau_1 \triangleright \tau_2), \Gamma' \vdash (\Sigma \vdash \tau) \triangleright \tau'$ **holds**, then $G[\alpha/\sigma][c/G'] \Rightarrow \Gamma, \Gamma'[\alpha/\sigma] \vdash (\Sigma[\alpha/\sigma] \vdash \tau[\alpha/\sigma]) \triangleright \tau'[\alpha/\sigma]$ **holds**.

- If $M \Rightarrow a : \Gamma, \alpha, (c : \tau_1 \triangleright \tau_2), \Gamma' \vdash \tau$ holds, then $M[\alpha/\sigma][c/G'] \Rightarrow a : \Gamma, \Gamma'[\alpha/\sigma] \vdash \tau[\alpha/\sigma]$ holds.

Proof. By induction using Lemma 55 for types and environments. The only interesting rule is COERVAR. There are two cases depending whether the coercion variable to instantiate is the same as the one of the judgment. If we have $c \Rightarrow \Gamma, \alpha, (c : \tau_1 \triangleright \tau_2), \Gamma' \vdash \tau_1 \triangleright \tau_2$ then we have to show that $G' \Rightarrow \Gamma, \Gamma'[\alpha/\sigma] \vdash \tau_1[\alpha/\sigma] \triangleright \tau_2[\alpha/\sigma]$ holds which we do by Lemma 54. If we have $c' \Rightarrow \Gamma, \alpha, (c : \tau_1 \triangleright \tau_2), \Gamma' \vdash \tau'_1 \triangleright \tau'_2$ then we have to show $c' \Rightarrow \Gamma, \Gamma'[\alpha/\sigma] \vdash \tau'_1 \triangleright \tau'_2$ which holds. \square

There are two soundness lemma: one for the explicit version and one for the implicit version. Both are done using a subject reduction and progress lemma. Subject reduction tells that if an term M is well-typed and reduces to N , then N is also well-typed with the same type and environment.

Lemma 59 (Explicit subject reduction). *If $M \Rightarrow a : \Gamma \vdash \tau$ and $M \rightsquigarrow_{\beta_L} N$ hold, then $N \Rightarrow b : \Gamma \vdash \tau$ holds where $b = \lfloor N \rfloor$.*

Proof. By induction on $M \rightsquigarrow_{\beta_L} N$. Rule REDCTX is done by cases on the context E . Rule REDAPP uses Lemma 57 and the fact that the drop function commutes with substitution. Rule REDWEAK uses Lemma 54 to show that M also has type τ under the well-formed extended environment Γ, Σ . Rules REDPOLY, REDPOLYL, and REDPOLYU use Lemma 58.

For rule REDETAARR, we have $G_1 \Rightarrow \Gamma, \Sigma \vdash \tau' \triangleright \tau$ and $G_2 \Rightarrow \Gamma \vdash (\Sigma \vdash \sigma) \triangleright \sigma'$ for the coercions and $M \Rightarrow a : \Gamma, \Sigma, (x : \tau) \vdash \sigma$ for the term. We rename x into a fresh term variable y in the right-hand side of the reduction. By Lemma 54 we have $G_2 \Rightarrow \Gamma, (y : \tau') \vdash (\Sigma \vdash \sigma) \triangleright \sigma'$. It remains to show that $M[x/G_1\langle y \rangle] \Rightarrow a[x/y] : \Gamma, (y : \tau'), \Sigma \vdash \sigma$ holds. In order to use Lemma 57, we need to show two subgoals. We first use Lemma 54 to show $M \Rightarrow a : \Gamma, (y : \tau'), \Sigma, (x : \tau) \vdash \sigma$. We then show $G_1\langle y \rangle \Rightarrow y : \Gamma, (y : \tau'), \Sigma \vdash \tau$ by Lemma 54 and rules TERMCOER and TERMVAR. Rules REDETAPROD is similar. \square

We classify values according to their types in the following lemma. This lemma, called explicit classification (canonical forms adapted to strong reduction), is used by the progress lemma, which follows. By looking at the type of a value, we get to know its form. This lemma is the natural extension of the canonical forms lemma to open terms, as our definition of values is the natural extension of values to open terms.

Lemma 60 (Explicit classification). *If $v \Rightarrow a : \Gamma \vdash \tau$ holds, then either v is a prevalue or the following assertions hold.*

- If τ is $\tau_1 \rightarrow \tau_2$, then v is $\lambda(x : \tau_1) v'$.
- If τ is $\tau_1 \times \tau_2$, then v is $\langle v_1, v_2 \rangle$.
- If τ is $\forall \alpha \tau'$, then v is $\Lambda \alpha \langle v' \rangle$.
- If τ is $\forall (\alpha \triangleright \sigma) \tau'$, then v is $\Lambda (\alpha \triangleright c : \sigma) \langle v' \rangle$.
- If τ is $\forall (\alpha \triangleleft \sigma) \tau'$, then v is $\Lambda (\alpha \triangleleft c : \sigma) \langle v' \rangle$.
- τ is not \perp .

Proof. By induction on v and inversion of $v \Rightarrow a : \Gamma \vdash \tau$. \square

The progress lemma tells that well-typed terms are either values or reducible. We make no distinction whether it is a β - or ι -reduction step.

Lemma 61 (Explicit progress). *If $M \Rightarrow a : \Gamma \vdash \tau$ holds, then either M is a value or there is N such that $M \rightsquigarrow_{\beta\iota} N$ holds.*

Proof. Let's assume that M is not a value. By induction on $M \Rightarrow a : \Gamma \vdash \tau$. Because of strong reduction, we can assume that all sub-terms of M are values, otherwise, M reduces using rule REDCTX. Rules TERMVAR, TERMLAM, and TERMPAIR are values. Rules TERMAPP, TERMFST, and TERMSND use Lemma 60 to reduce by rules REDAPP, REDFST, and REDSND, respectively.

For rule TERMCOER, we proceed by cases on $G \Rightarrow \Gamma \vdash (\Sigma \vdash \tau) \triangleright \sigma$. Rules COERREFL, COERTRANS, and COERWEAK reduce by rules REDREFL, REDTRANS, and REDWEAK. Rules COERVAR, COERTLAM, COERTOP, COERTLAML, and COERTLAMU are values. Rules COERTAPP, COERETAARR, COERETAPROD, COERTAPPL, and COERTAPPU use Lemma 60 to reduce by rules REDPOLY, REDETAARR, REDEETAPROD, REDPOLYL, and REDPOLYU, respectively.

The remaining rule is COERBOT which we refute with Lemma 60: there is no value of type \perp . \square

Subject reduction and progress together proves the soundness lemma. An explicit term is sound if it reduces to a value.

Proposition 62 (Explicit soundness). *If $M \Rightarrow a : \Gamma \vdash \tau$ holds, then M terminates to a value.*

Proof. By Lemma 48 and 50, M terminates to N . By Lemma 59 we have $N \Rightarrow [N] : \Gamma \vdash \tau$. And by Lemma 61 we have that N is a value. \square

The same schema for soundness works for the implicit version. We prove subject reduction and progress for the implicit version from the explicit version using the bisimulation property.

Lemma 63 (Implicit subject reduction). *If $a : \Gamma \vdash \tau$ and $a \rightsquigarrow b$ hold, then $b : \Gamma \vdash \tau$ holds.*

Proof. By Lemma 45 we have $M \Rightarrow a : \Gamma \vdash \tau$ for some M . By Lemma 53 we have $M \rightsquigarrow_{\iota}^* \rightsquigarrow_{\beta} N$ such that $[N] = b$. By Lemma 59 we have $N \Rightarrow b : \Gamma \vdash \tau$. We conclude with Lemma 45. \square

The progress lemma for the implicit version of the type system resemble its explicit version. However, the proof may be short-cut using the bisimulation property and the explicit progress lemma.

Lemma 64 (Implicit Progress). *If $a : \Gamma \vdash \tau$ holds, then a is a value or there is a term b such that $a \rightsquigarrow b$ holds.*

Proof. Let's assume a is not a value. By Lemma 45 we have $M \Rightarrow a : \Gamma \vdash \tau$ for some M . We proceed by induction on the strong normalization of M (Lemma 48). If M is a value, then a is a value by Lemma 52 and we are done. If M is not a value, then by Lemma 61 we have $M \rightsquigarrow_{\beta\iota} N$. If the reduction is a ι -step, we call the induction hypothesis since $[N] = [M] = a$. If the reduction is a β -step, then by Lemma 52, we have $a \rightsquigarrow [N]$. \square

Now that we have shown explicit subject reduction, we can show the strong normalization of the implicit reduction. This result is useful to prove the soundness of the implicit language.

Proposition 65 (Implicit termination). *If $a : \Gamma \vdash \tau$ holds, then a strongly normalizes.*

Proof. By Lemma 45 we have $M \Rightarrow a : \Gamma \vdash \tau$ with $\llbracket M \rrbracket = a$. We proceed by induction on the strong normalization of M . If $a \rightsquigarrow b$, by Lemma 53 we have $M \rightsquigarrow_{\iota}^* \rightsquigarrow_{\beta} N$ with $\llbracket N \rrbracket = b$. By Lemma 59, we have $N \Rightarrow x : \Gamma \vdash \tau$, which allows us to call the induction hypothesis to conclude. \square

Similarly to the explicit soundness lemma, the implicit soundness lemma tells that well-typed lambda-terms strongly normalize to a value.

Proposition 66 (Implicit soundness). *If $a : \Gamma \vdash \tau$ holds, then a terminates to a value.*

Proof. By Lemma 65, a terminates to b . By Lemma 63 we have $b : \Gamma \vdash \tau$. And by Lemma 64 we have that b is a value. \square

4.5 Expressivity

To prove that a type system T_2 is at least as expressive than a type system T_1 , we need to show that if a term a is well-typed in T_1 then it is also well-typed in T_2 . This implies that there are more well-typed terms in the most expressive type-system than the less expressive one. Keeping in mind that well-typed terms are sound and that type systems are conservative, an expressive type-system rejects fewer sound terms.

It is sufficient to show the inclusion between the implicit versions since explicit and implicit versions are equivalent. This simplifies the notations, but the proofs are mainly the same.

4.5.1 System F

In this section, we show that System F is equivalent to the base system extended with polymorphism. For each judgment of System F, we show that the equivalent judgment in System F_{ι}^p also holds. Actually, we only need the polymorphism feature on top of the base system to include System F. And reciprocally, the base system extended with polymorphism is included into System F. As a consequence System F is equivalent to polymorphism, which is what is expected.

Lemma 67. *We write $(J)_s$ for judgments in System F (as source language), see Section 3.2, and $(J)_t$ for those in the base system extended with the polymorphism extension (as target language). The following assertions hold.*

- *If $(\Gamma \text{ env})_s$ holds, then $(\Gamma \text{ env})_t$ holds, and reciprocally.*
- *If $(\Gamma \vdash \tau \text{ type})_s$ holds, then $(\Gamma \vdash \tau \text{ type})_t$ holds, and reciprocally.*
- *If $(\Gamma \vdash a : \tau)_s$ holds, then $(a : \Gamma \vdash \tau)_t$ holds, and reciprocally.*
- *If $(\Gamma \vdash (\Sigma \vdash \tau) \triangleright \sigma)_t$ and $(\Gamma, \Sigma \vdash a : \tau)_s$ hold, then $(\Gamma \vdash a : \sigma)_s$ holds.*

Since there is no counter-part of coercions in System F, the last assertion tells that coercions are translated to term derivations transformation.

Proof. First, we prove the implication from source to target and then we prove the reciprocal from target to source. Notice that the last assertion only occurs for the reciprocal. For both directions we proceed by mutual induction on all judgments.

For the implication, environment and type well-formedness rules are trivial since they are exactly the same. The same argument holds for the first six term judgment rules (those of the STLC): `TERMVAR`, `TERMLAM`, `TERMAPP`, `TERMPAIR`, `TERMFST`, and `TERMSND`. The two interesting cases are rules `TERMGEN` and `TERMINST`.

For rule `TERMGEN`, we have $(\Gamma, \alpha \vdash a : \tau)_s$ and we need to prove $(a : \Gamma \vdash \forall \alpha \tau)_t$. By induction hypothesis we have $(a : \Gamma, \alpha \vdash \tau)_t$. We use rule `TERMCOR` on the preceding derivation. It remains to prove $(\Gamma \vdash (\alpha \vdash \tau) \triangleright \forall \alpha \tau)_t$ which we do with rule `COERTLAM`.

For rule `TERMINST`, we have $(\Gamma \vdash a : \forall \alpha \tau)_s$ and $(\Gamma \vdash \sigma \text{ type})_s$ and we need to prove $(a : \Gamma \vdash \tau[\alpha/\sigma])_t$. By induction hypothesis we have $(a : \Gamma \vdash \forall \alpha \tau)_t$ and $(\Gamma \vdash \sigma \text{ type})_t$. We use rule `TERMCOR` on the preceding term derivation. It remains to prove $(\Gamma \vdash \forall \alpha \tau \triangleright \tau[\alpha/\sigma])_t$ which we do with rule `COERTAPP`.

Let's now prove the reciprocal. Environment and type well-formedness rules are trivial again since they are exactly the same (notice that we only consider the base system extended with the polymorphism extension). The same argument holds for the first six term judgment rules (those of the STLC): `TERMVAR`, `TERMLAM`, `TERMAPP`, `TERMPAIR`, `TERMFST`, and `TERMSND`. For rule `TERMCOR`, we use the last assertion on the induction hypothesis.

It remains to prove the last assertion. Rule `COERREFL` is done by hypothesis. Rule `COERTTRANS` is done by composition. Rule `COERVAR` cannot occur since by extraction we have $(\Gamma \text{ env})_s$ which implies that the environment Γ do not contain any coercion variables. Rule `COERWEAK` uses the weakening lemma. Finally the two interesting cases are rules `COERTLAM` and `COERTAPP`.

For rule `COERTLAM` we use rule `TERMGEN` and for rule `COERTAPP` we use rule `TERMINST`. □

4.5.2 System F_η

In this section, we show that System F_η is equivalent to the base system extended with polymorphism and η -expansion. For each judgment of System F_η , we show that the equivalent judgment in System F_t^p also holds. Actually, we only need the polymorphism and η -expansion features on top of the base system to include F_η . And reciprocally, the base system extended with polymorphism and η -expansion is included into F_η . As a consequence F_η is equivalent to polymorphism and η -expansion, or equivalently extends System F with η -expansion as its name suggests. We write $(J)_s$ for judgments in System F_η and $(J)_t$ for those in the base system extended with the polymorphism and the η -expansion extensions. We define $\forall \Sigma \tau$ inductively on Σ .

$$\begin{aligned} \forall \emptyset \tau &\stackrel{\text{def}}{=} \tau \\ \forall (\Sigma, \alpha) \tau &\stackrel{\text{def}}{=} \forall \Sigma (\forall \alpha \tau) \end{aligned}$$

Lemma 68. *The following assertions hold.*

- If $(\Gamma \text{ env})_s$ holds, then $(\Gamma \text{ env})_t$ holds, and reciprocally.
- If $(\Gamma \vdash \tau \text{ type})_s$ holds, then $(\Gamma \vdash \tau \text{ type})_t$ holds, and reciprocally.
- If $(\Gamma \vdash a : \tau)_s$ holds, then $(a : \Gamma \vdash \tau)_t$ holds, and reciprocally.
- If $(\Gamma \vdash \tau \triangleright \sigma)_s$ holds, then $(\Gamma \vdash \tau \triangleright \sigma)_t$ holds.

- If $(\Gamma \vdash (\Sigma \vdash \tau) \triangleright \sigma)_t$ holds, then $(\Gamma \vdash \forall \Sigma \tau \triangleright \sigma)_s$ holds.

Since there is no counter-part of coercions in System F_η , but just a notion of type containment (that do not extend the environment), the last assertion tells that coercions are translated to type containment where the source type is the type version of the source typing of the coercion.

Proof. First we prove the implication from source to target and then we prove the reciprocal from target to source. Notice that the last assertion only occurs for the reciprocal. For both directions we proceed by mutual induction on all judgments.

For the implication, environment and type well-formedness rules are trivial since they are exactly the same. The same argument holds for the first six term judgment rules (those of the STLC): `TERMVAR`, `TERMLAM`, `TERMAPP`, `TERMPAIR`, `TERMFST`, and `TERMSND`. The next two cases, namely rules `TERMGEN` and `TERMINST`, are similar to the proof for System F inclusion. Finally, for rule `TERMCONT`, we use rule `TERMCOER`.

It remains to show the implication for the containment judgment. Rule `CONTREFL` uses rule `COERREFL` and rule `CONTTRANS` uses rule `COERTRANS` and weakening. Rule `CONTTLAM` uses rule `COERTLAM`. Rule `CONTTAPP` uses rule `COERTAPP`. Rule `CONTARR` uses rule `COERETAARR` and weakening. Rule `CONTPROD` uses rule `COERETAPROD`.

For rule `CONTCONGR`, we have $(\Gamma, \alpha \vdash \tau \triangleright \sigma)_t$ by induction hypothesis and we show $(\Gamma \vdash \forall \alpha \tau \triangleright \forall \alpha \sigma)_t$. At first we apply rule `COERWEAK` which leaves us to show $(\Gamma \vdash (\alpha \vdash \forall \alpha \tau) \triangleright \forall \alpha \sigma)_t$. We apply a first time rule `COERTRANS` with rule `COERTLAM` which leaves us with $(\Gamma, \alpha \vdash \forall \alpha \tau \triangleright \sigma)_t$ to prove. We apply a second time rule `COERTRANS` with the induction hypothesis this time. This leaves us $(\Gamma, \alpha \vdash \forall \alpha \tau \triangleright \tau)_t$ to prove, for which we use rule `COERTAPP` with the type argument α .

For rule `CONTDISTARR`, we take α fresh. After the use of rule `COERWEAK`, it remains to show that $(\Gamma \vdash (\alpha \vdash \forall \alpha \tau \rightarrow \sigma) \triangleright \tau \rightarrow \forall \alpha \sigma)_t$ holds. We apply rule `COERTRANS` on $(\Gamma \vdash (\alpha \vdash \tau \rightarrow \sigma) \triangleright \tau \rightarrow \forall \alpha \sigma)_t$ and $(\Gamma, \alpha \vdash \forall \alpha \tau \rightarrow \sigma \triangleright \tau \rightarrow \sigma)_t$. We prove the last one with rule `COERTAPP` with type argument α . We prove the first one with rule `COERETAARR`, which gives one goal $(\Gamma, \alpha \vdash \tau \triangleright \tau)_t$ for the argument type and $(\Gamma \vdash (\alpha \vdash \sigma) \triangleright \forall \alpha \sigma)_t$ for the return type. We prove the first one with reflexivity and the second one with rule `COERTLAM`. The coercion witness for `dist` in System F_t^P is thus $*((\diamond \xrightarrow{\tau} \Lambda \alpha) \circ \cdot \alpha)$.

Rule `CONTDISTPROD` is similar. After rule `COERWEAK`, we show $(\Gamma \vdash (\alpha \vdash \forall \alpha \tau \times \sigma) \triangleright \forall \alpha \tau \times \forall \alpha \sigma)_t$ with α a fresh type variable. We first instantiate with α then use the product η -expansion and type abstraction twice.

Let's now prove the reciprocal. Environment and type well-formedness rules are trivial again since they are exactly the same. The same argument holds for the first six term judgment rules (those of the STLC): `TERMVAR`, `TERMLAM`, `TERMAPP`, `TERMPAIR`, `TERMFST`, and `TERMSND`. For rule `TERMCOER`, we use the last assertion on the induction hypothesis.

It remains to prove the last assertion. Rule `COERREFL` is done with rule `CONTREFL`. For rule `COERTRANS`, we use rule `CONTTRANS` on $(\Gamma \vdash \forall \Sigma_2, \Sigma_1 \tau_1 \triangleright \forall \Sigma_2 \tau_2)_s$ and $(\Gamma \vdash \forall \Sigma_2 \tau_2 \triangleright \tau_3)_s$. The last one holds by induction hypothesis. For the first one we use rule `CONTCONGR` repeatedly to get $(\Gamma, \Sigma_2 \vdash \forall \Sigma_1 \tau_1 \triangleright \tau_2)_s$ which holds by induction hypothesis. Rule `COERVAR` cannot occur since by extraction we have $(\Gamma \text{ env})_s$ which implies that the environment Γ does not contain any coercion variables. For rule `COERWEAK` we use repeatedly a combination of rules `CONTTRANS`, `CONTCONGR`, and `CONTTLAM` to show $(\Gamma \vdash \tau \triangleright \forall \Sigma \tau)_s$. Rule `COERTLAM` uses rule `CONTREFL`. `COERTAPP` uses rule `CONTTAPP`. The last remaining rules are `COERETAARR` and `COERETAPROD`.

For rule COERETAARR, we use rule CONTTRANS on $(\Gamma \vdash \forall \Sigma (\tau \rightarrow \sigma) \triangleright \forall \Sigma (\tau' \rightarrow \sigma))_s$ and $(\Gamma \vdash \forall \Sigma (\tau' \rightarrow \sigma) \triangleright \tau' \rightarrow \sigma')_s$. For the first containment we use rule CONTCONGR repeatedly to get $(\Gamma, \Sigma \vdash \tau \rightarrow \sigma \triangleright \tau' \rightarrow \sigma)_s$ which we show with rule CONTARR and the induction hypothesis. For the second containment we use rule CONTTRANS on $(\Gamma \vdash \forall \Sigma (\tau' \rightarrow \sigma) \triangleright \tau' \rightarrow \forall \Sigma \sigma)_s$ and $(\Gamma \vdash \tau' \rightarrow \forall \Sigma \sigma \triangleright \tau' \rightarrow \sigma')_s$. For the first containment we use rules CONTTRANS, CONTCONGR, and CONTDISTARR. For the second containment we use rules CONTARR with the induction hypothesis. Rule COERETAPROD is similar. \square

4.5.3 MLF

In this section, we show that MLF is included in the base system extended with lower bounded polymorphism and the bottom type. For each judgment of MLF, we show that the equivalent judgment in System F_t^p also holds. Actually, we only need the lower bounded polymorphism and bottom type feature on top of the base system to include MLF. And the reciprocal is almost true: MLF would include the base system extended with lower bounded polymorphism and bottom type, if recursive bounds were present in MLF. In other words, MLF is equivalent to the base system extended with non-recursive lower bounded polymorphism and the bottom type. We write $(J)_s$ for judgments in MLF and $(J)_t$ for those in the base system extended with lower bounded polymorphism and bottom type.

Lemma 69. *The following assertions hold.*

- *If $(\Gamma \text{ env})_s$ holds, then $(\Gamma \text{ env})_t$ holds.*
- *If $(\Gamma \vdash \tau \text{ type})_s$ holds, then $(\Gamma \vdash \tau \text{ type})_t$ holds.*
- *If $(\Gamma \vdash \tau \triangleright \sigma)_s$ holds, then $(\Gamma \vdash \tau \triangleright \sigma)_t$ holds.*
- *If $(\Gamma \vdash a : \tau)_s$ holds, then $(a : \Gamma \vdash \tau)_t$ holds.*

Proof. We proceed by mutual induction on all judgments. Environment and type well-formedness rules are trivial since they are exactly the same. The same argument holds for the first six term judgment rules (those of the STLC): TERMVAR, TERMLAM, TERMAPP, TERMPAIR, TERMFST, and TERMSND. Rule TERMTABS is a conjunction of rules COERTLAM and TERMCOER. Rule TERMTAPP uses rule TERMCOER.

For the instantiation judgment, rule INSTBOT uses rule COERBOT. Rule INSTID is rule COERREFL. And rule INSTCOMP is rule COERTRANS. Rule INSTABSTR is rule COERVAR. For rule INSTINTRO, we use rule COERWEAK after rule COERTLAM since the abstract type variable is not used. Rule INSTELIM uses rule COERTAPPL with the bound for the argument type and the reflexivity coercion (rule COERREFL) for the argument coercion.

For rule INSTUNDER, we use rule COERWEAK on $\Gamma \vdash (\alpha, (\tau \triangleright \alpha) \vdash \forall (\alpha \triangleleft \tau) \sigma_1) \triangleright \forall (\alpha \triangleleft \tau) \sigma_2$, which we prove by using rule COERTRANS twice. First, we prove $\Gamma, \alpha, (\tau \triangleright \alpha) \vdash \forall (\alpha \triangleleft \tau) \sigma_1 \triangleright \sigma_1$ by rule COERTAPPL with α as the type argument and c as the coercion argument. Then, we prove $\Gamma, \alpha, (\tau \triangleright \alpha) \vdash \sigma_1 \triangleright \sigma_2$ by induction hypothesis. And we finally prove $\Gamma \vdash (\alpha, (\tau \triangleright \alpha) \vdash \sigma_2) \triangleright \forall (\alpha \triangleleft \tau) \sigma_2$ by rule COERTLAM.

For rule INSTINSIDE, we use rule COERWEAK on $\Gamma \vdash (\alpha, (\tau_2 \triangleright \alpha) \vdash \forall (\alpha \triangleleft \tau_1) \sigma) \triangleright \forall (\alpha \triangleleft \tau_2) \sigma$, which we prove by using rule COERTRANS. First we prove $\Gamma, \alpha, (\tau_2 \triangleright \alpha) \vdash \forall (\alpha \triangleleft \tau_1) \sigma \triangleright \sigma$ by rule COERTAPPL with α as the type argument. For the coercion argument $\Gamma, \alpha, (\tau_2 \triangleright \alpha) \vdash \tau_1 \triangleright \alpha$, we use rule COERTRANS with the induction hypothesis and c . Then we prove $\Gamma \vdash (\alpha, (\tau_2 \triangleright \alpha) \vdash \sigma) \triangleright \forall (\alpha \triangleleft \tau_2) \sigma$ by rule COERTLAM. \square

We conclude from this inclusion, that MLF is sound and strongly normalizing. The proof of strong normalization of MLF is actually one of the contribution of this thesis.

4.5.4 System $F_{<}$.

We remind that the version of System $F_{<}$ we describe in Section 3.6 corresponds to the most expressive version of System $F_{<}$, namely System F-bounded [5]. For each judgment of System $F_{<}$, we show that the equivalent judgment in System F_t^p also holds. Actually, we only need the upper bounded polymorphism, the top type, and η -expansion features on top of the base system to include $F_{<}$. However, unlike for System F and System F_η , and like for MLF, the reciprocal is not completely true. The base system extended with upper bounded polymorphism, top, and η -expansion is strictly more expressive than System $F_{<}$. We explain why after the proof. We write $(J)_s$ for judgments in System $F_{<}$ and $(J)_t$ for those in the base system extended with upper bounded polymorphism, the top type, and η -expansion.

Lemma 70. *The following assertions hold.*

- *If $(\Gamma \text{ env})_s$ holds, then $(\Gamma \text{ env})_t$ holds.*
- *If $(\Gamma \vdash \tau \text{ type})_s$ holds, then $(\Gamma \vdash \tau \text{ type})_t$ holds.*
- *If $(\Gamma \vdash \tau \triangleright \sigma)_s$ holds, then $(\Gamma \vdash \tau \triangleright \sigma)_t$ holds.*
- *If $(\Gamma \vdash a : \tau)_s$ holds, then $(a : \Gamma \vdash \tau)_t$ holds.*

Proof. We proceed by mutual induction on all judgments. Environment and type well-formedness rules are trivial since they are exactly the same. The same argument holds for the first six term judgment rules (those of the STLC): TERMVAR, TERMLAM, TERMAPP, TERMPAIR, TERMFST, and TERMSND. Rule TERMGEN is a conjunction of rules COERTLAMU and TERMCOER. Rule TERMINST is a conjunction of rules COERTAPPU and TERMCOER. Rule TERMSUB uses rule TERMCOER.

For the subtyping judgment, rules SUBREFL, SUBTRANS, SUBVAR, SUBTOP, SUBARR, and SUBPROD use rules COERREFL, COERTRANS, COERVAR, COERTOP, COERETAARR, and COERETAPROD respectively. For rule SUBCONGR, we use rule COERWEAK on $\Gamma \vdash (\alpha, (\alpha \triangleright \tau') \vdash \forall(\alpha \triangleright \tau) \sigma) \triangleright \forall(\alpha \triangleright \tau') \sigma'$ which we prove with two consecutive uses of rule COERTRANS. We first prove $\Gamma, \alpha, (\alpha \triangleright \tau') \vdash \forall(\alpha \triangleright \tau) \sigma \triangleright \sigma$ with rule COERTAPPU with α for the type argument and the induction hypothesis $\Gamma, \alpha, (\alpha \triangleright \tau') \vdash \alpha \triangleright \tau$ for the coercion argument. Then we prove $\Gamma, \alpha, (\alpha \triangleright \tau') \vdash \sigma \triangleright \sigma'$ with the second induction hypothesis. And we finally prove $\Gamma \vdash (\alpha, (\alpha \triangleright \tau') \vdash \sigma') \triangleright \forall(\alpha \triangleright \tau') \sigma'$ with rule COERTLAMU. \square

The reciprocal does not hold because System $F_{<}$ misses a distributivity rule like we can find in System F_η . Distributivity naturally comes from polymorphism and η -expansion when expressed as coercions. System F_η and System $F_{<}$ resemble each other. They are composed of two type system features: η -expansion and some sort of polymorphism. However, they do not handle polymorphism in the same manner. In System F_η , polymorphism is in the coercion judgment. While in System $F_{<}$, (upper bounded) polymorphism is only available in the term judgment and thus not composable with the other type system features, in our case η -expansion (called subtyping in System $F_{<}$). More concretely, upper bounded instantiation is not a subtyping rule, but only a typing rule. While in System F_η , type instantiation is a containment rule.

The distributivity rule is not the only rule derivable in System $F_{<}$: if upper bounded polymorphism was in the coercion judgment. The congruence rule for upper bounded polymorphism would also be derivable. A general statement about erasable type system features (like polymorphism, recursive types, coercion abstraction) is that they come with an introduction and elimination coercion rule, from which we derive their congruence rule and distributivity rule for those about abstraction.

4.5.5 Summary

We can summarize all these inclusions and equivalences into the following table.

Type system features	F	F_η	MLF	$\supseteq F_{<}$	F_t^p
Polymorphism	✓	✓	-	-	✓
Eta-expansion	-	✓	-	✓	✓
Bottom	-	-	✓	-	✓
Top	-	-	-	✓	✓
Lower bounded polymorphism	-	-	✓	-	✓
Upper bounded polymorphism	-	-	-	✓	✓

System F is equivalent to the base system extended with polymorphism. System F_η is equivalent to the base system extended with polymorphism and η -expansion. MLF is equivalent to the base system extended with the bottom type and lower bounded polymorphism. Actually, the equivalence is almost true, but for recursive bounds. System $F_{<}$ is included in the base system extended with η -expansion, the top type, and upper bounded polymorphism. System $F_{<}$, which corresponds to the most expressive version of System $F_{<}$: [5], is in fact strictly included into this extension because the distributivity rule is derivable in the extended system but not in System $F_{<}$. In this sense, System $F_{<}$ is not complete according to the features it presents, whereas System F and System F_η are complete with respect to their features. The reason why the distributivity rule is not derivable is that the upper bounded polymorphism feature is not a subtyping (or coercion) rule in $F_{<}$. It is only a typing rule at the term level. The subtyping judgment only contains the congruence rule for upper bounded polymorphism. Finally, System F_t^p contains all features and thus includes all presented type systems.

4.6 Beyond parametric coercion abstraction

In this chapter, we have defined a coercion type system where all features are expressed as coercions and thus fully composable and erasable. We called this coercion type system System F_t^p . We defined it as a base system corresponding to the STLC and a series of orthogonal and composable features. We showed that the STLC, System F , System F_η , MLF, and System $F_{<}$ can be seen as the base system with some additional features. Notice however that System $F_{<}$ is not as expressive as it naturally gets in this approach. As a corollary, all these type systems are included in System F_t^p which contains all features.

However, System F_t^p is not as expressive as we might want. It has some restrictions and as a consequence it does not contain **Constraint ML**, even though polymorphism is not even first class in **Constraint ML**. The cause of this restriction is coercion abstraction. As a matter of

fact, upper and lower bounded polymorphism are forms of coercion abstraction, but they have the particularity that they only abstract over parametric coercions: coercions which are either from an abstract type or to an abstract type. In order to describe the feature of **Constraint ML** we would need an unrestricted coercion abstraction of the form $\Lambda(c : \tau \triangleright \sigma)$. We discuss this point in Section 4.6.1. Conclusions of this section leads us to a discussion about push in Section 4.6.2.

4.6.1 Unrestricted coercion abstraction

We could define a type for coercion polymorphism written $(\tau \triangleright \sigma) \Rightarrow \rho$ abstracting over the coercion $\tau \triangleright \sigma$ with body ρ . We would then naturally expect the following associated typing rules:

$$\begin{array}{c} \text{CoERCLAM} \\ \Lambda(c : \tau \triangleright \sigma) \Rightarrow \Gamma \vdash ((c : \tau \triangleright \sigma) \vdash \rho) \triangleright (\tau \triangleright \sigma) \Rightarrow \rho \end{array} \qquad \begin{array}{c} \text{CoERCAPP} \\ \frac{G \Rightarrow \Gamma \vdash \tau \triangleright \sigma}{\bullet G \Rightarrow \Gamma \vdash (\tau \triangleright \sigma) \Rightarrow \rho \triangleright \rho} \end{array}$$

The problem is that in doing so (without any additional restriction) we break backward simulation. To see why, let's consider the following explicit term where c is a coercion variable: $(c(\lambda(x : \tau) M)) N$. It drops to $(\lambda x [M]) [N]$ which is a redex, and thus can reduce. However, without allowing decomposition of abstract coercions (called push and described in Section 4.6.2), the explicit term is stuck and cannot do the same step as its dropped lambda term.

With push, backward simulation seems to be restored, but this breaks implicit progress and termination. Let's consider this non-terminating term $(\Lambda(c_1 : \alpha \triangleright \alpha \rightarrow \alpha) \circ \Lambda(c_2 : \alpha \rightarrow \alpha \triangleright \alpha)) \langle a (c_2 \langle a \rangle) \rangle$ where $a = \lambda(x : \alpha) c_1 \langle x \rangle x$. This is the usual omega term. Its erasure loops, while this term builds always growing coercions by projection.

Independently of a push reduction rule, when abstracting over coercions, we need to make sure that we only abstract over consistent coercions. We shouldn't allow to abstract over c of type $\tau \times \sigma \triangleright \tau \rightarrow \sigma$. Let's consider the term $c \langle \langle M, N \rangle \rangle M$ which is well-typed. It's erasure is $\langle [M], [N] \rangle [M]$ which is stuck therefore breaking the progress lemma.

A solution to avoid considering consistency is to keep coercions during reduction. The reduction does not go wrong, but coercions are no more erasable. GHC is doing something related for their first-class coercions. They only erase the content of coercions but not their existence. Besides, they don't reduce under coercion abstraction. Coercions can thus be represented with 0-bit, *i.e.* the unit term. We extend our approach with 0-bit coercions in Section 6.1.5.

GHC adopts another solution for their top-level coercions. They have a consistency lemma which is needed to prove progress for the implicit language. In our case this lemma is ν -classification (Lemma 51). In GHC, they have a more restrictive lemma that asks all coercions not to touch their head type constructor. Our lemma is more subtle since it only ask this condition for non-erasable type constructors.

The solution used in System F_l^p is inspired from $F_{<}$ and MLF. It allows coercion abstraction only for parametric coercion types. In $F_{<}$, the parametricity is on the argument type and in MLF it is on the return type. In System F_l^p we allow both. This restriction prevents to build a coercion between two arrow types using a coercion variable. This is proved in the classification lemma (Lemma 51) which is needed to prove the bisimulation property.

The solution we develop in Chapter 5 with System F_{cc} is to ask all abstract coercions to be syntactically inhabited, which implies consistency. The minimum requirement for soundness would be to ask a semantic witness that the abstract coercion is inhabited. This extension is discussed in Section 6.1.10.

4.6.2 Push

We call *abstract*, coercions containing coercion variables in places such that reduction cannot progress with the reduction rules of System F_t^P .

In our current setting, only concrete coercions are decomposed. For example, an arrow η -expansion coercion applied to a function can be reduced (with Rule `REDETAARR`) into a new function that coerces its argument and pass it to the previous function and then coerces the result. Decomposing abstract coercions is not needed, because reduction can always bring concrete coercions in places where decomposition of coercions would be needed. However, if we relax our restriction on coercion abstraction, we may have abstract coercions in-between redexes blocking the reduction that could be performed if the coercion were erased—thus breaking backward simulation. In this case, backward simulation may be recovered by adding a decomposition rule for abstract coercions—and associated coercion constructs.

Interestingly, decomposing coercions makes the backward simulation much easier, since wedges can now be decomposed instead of blocking the evaluation. However, the decomposition introduces new operations on abstract coercions that must eventually be reduced when coercions become concrete: proof obligations are just moved elsewhere. Currently, explicit progress, consistency, and backward simulation are proved using coercion values in ι -classification. In order to have coercion values with push, we need to add reduction rules for coercions.

Figure 4.18 shows how a coercion in-between a redex can be decomposed. The initial term $G\langle\lambda(x : \tau_1) M\rangle N$, on the left, is the application of the function $\lambda(x : \tau_1) M$ coerced by G to the argument N . Its erasure is the application of the erasure of $\lambda(x : \tau_1) M$, namely $\lambda x a$, to the erasure of N , namely b , that is, the redex $(\lambda x a) b$. Can we make the redex apparent on the source term, without inspecting the coercion G , *i.e.* in a way that may also work when G is an abstract coercion (a coercion variable for instance)?

The terms are drawn upside-down to mimic their typing derivations (terms are explicitly typed and thus isomorphic to their typing derivations). Boxes represent scopes. The coercion G can extend the environment from Γ_2 to Γ_1 with the erasable environment Σ_1 ($\Gamma_1 = \Gamma_2, \Sigma_1$). The blue (dark gray for black and white prints) box delimits the separation between the Γ_2 environment (which is outside with a white background color) and the Γ_1 environment (which is inside with a blue background color). This blue part is more precisely the application of the environment action Σ_1 to Γ_2 (see Section 6.1.11). The lambda constructor opens a scope and thus also changes the environment. Its effect is to add a term binding. We use a green (or light gray for black and white versions) background color to show when x is added to the environment. We add a tick to the environment meta-variable when it is extended with the binding for x (see the first two definitions on Figure 4.18).

The coercion G lies on the left-hand side of the application and outside of the abstraction. It can be moved on other branches of the application node (upward arrow) or inside the abstraction (downward arrow). The argument is coerced directly in the former case, but indirectly on each use of the function parameter in the latter case.

In both cases we need to decompose G of type $\Gamma_2 \vdash (\Sigma_1 \vdash \tau_1 \rightarrow \sigma_1) \triangleright \tau_2 \rightarrow \sigma_2$ into RG




Figure 4.18: Push

of type $\Gamma_2 \vdash (\Sigma_1 \vdash \sigma_1) \triangleright \sigma_2$ and LG of type $\Gamma_1 \vdash (\Sigma_2 \vdash \tau_2) \triangleright \tau_1$ (modulo the presence of x when we push under the lambda) where Σ_2 reverses the action of Σ_1 (see Section 6.1.11). If the only actions are bindings, like in System F_t^p and usual type systems, LG would have type $\Gamma_1 \vdash \tau_2 \triangleright \tau_1$, binding nothing, and the blue boxes would spread upward the LG node and not only downward.

Besides, for both reduction rules, the scope of G moves. It starts at RG and ends at LG (or it ends at the leaves of the typing derivation if we don't use environment actions). We can actually see this in the condition that $\Gamma_2, \Sigma_1, \Sigma_2$ (the environment once we leave LG) should behave as Γ_2 (the environment before we enter RG).

Fortunately, the two possible reductions are locally confluent: both terms reduce with a β -reduction on the rightmost term where N got substituted in M and both N and M are coerced with the left and right projections of G —and the whole term remains well-typed. Therefore, only one of either rule may be retained, indifferently, or both rules may be offered, simultaneously. It is also possible to add a unique rule doing both steps at the same time. This rule would then be labeled as a β -reduction rule since its erasure does one step of computation.

At this stage, we may consider the following typing rules and reduction rule:

$$\begin{array}{c}
\text{COERPUSHRIGHT} \\
\frac{G \Rightarrow \Gamma_2 \vdash (\Sigma_1 \vdash \tau_1 \rightarrow \sigma_1) \triangleright \tau_2 \rightarrow \sigma_2}{RG \Rightarrow \Gamma_2 \vdash (\Sigma_1 \vdash \sigma_1) \triangleright \sigma_2}
\end{array}
\qquad
\begin{array}{c}
\text{COERPUSHLEFT} \\
\frac{G \Rightarrow \Gamma_2 \vdash (\Sigma_1 \vdash \tau_1 \rightarrow \sigma_1) \triangleright \tau_2 \rightarrow \sigma_2}{LG \Rightarrow \Gamma_2 \vdash \tau_2 \triangleright \tau_1}
\end{array}$$

$$\begin{array}{c}
\text{REDPUSHARR} \\
G \langle \lambda(x : \tau) M \rangle N \rightsquigarrow_{\beta} RG \langle M[x/LG \langle N \rangle] \rangle
\end{array}$$

Still, we need further reduction rules to recover confluence of the language and in particular have a progress lemma for coercions, because we need coercion values for the ι -classification lemma and thus implicit progress. Indeed, the decomposition introduced new left and right destructors for coercions, which must eventually be reduced in order to show that the implicit language is sound.

We define a γ -reduction written $G \rightsquigarrow_{\gamma} G$ for coercions (see Section 6.1.12). The goal of this reduction relation is to give closed coercions a value which has to be a concrete coercion of System F_t^P . Some examples for this reduction relation are:

$$\begin{array}{ll}
R(G_1 \xrightarrow{\tau} G_2) \rightsquigarrow_{\gamma} G_2 & R\Diamond \rightsquigarrow_{\gamma} \Diamond \\
L(G_1 \xrightarrow{\tau} G_2) \rightsquigarrow_{\gamma} G_1 & L\Diamond \rightsquigarrow_{\gamma} \Diamond
\end{array}$$

It may be actually quite hard to find a complete set of rules, because we need to define the left and right destructors for all previous coercions. Moreover, this push mechanism complicates the framework for future coercion extensions.

One example of a difficult γ -reduction rule is when we want to reduce the left or right projections of $\cdot \rho \circ \Lambda \alpha$, which happens when we instantiate and apply an argument to a polymorphic function. This coercion has type $\Gamma \vdash (\alpha \vdash \tau \rightarrow \sigma) \triangleright \tau[\alpha/\rho] \rightarrow \sigma[\alpha/\rho]$. We want to extract two coercions of type $\Gamma \vdash (\alpha \vdash \sigma) \triangleright \sigma[\alpha/\rho]$ for the right projection and $\Gamma, \alpha \vdash \tau[\alpha/\rho] \triangleright \tau$ for the left projection. We recognize type instantiation for the right projection and the same coercion $\cdot \rho \circ \Lambda \alpha$ works. For the left projection however the situation is more complicated. We need the reverse operation of type instantiation.

One solution to this problem is to preserve the link between the right and left projections and factor the initial coercion. One way to do so is to have the right projection bind the left projection. The right projection would then be $\pi(c)G$ and the left would be c . We would then have the following coercion and reduction rules:

$$\begin{array}{c}
\text{COERPUSH} \\
\frac{G \Rightarrow \Gamma_2 \vdash (\Sigma_1 \vdash \tau_1 \rightarrow \sigma_1) \triangleright \tau_2 \rightarrow \sigma_2}{\pi(c)G \Rightarrow \Gamma_2 \vdash (\Sigma_1, (c : \tau_2 \triangleright \tau_1) \vdash \sigma_1) \triangleright \sigma_2}
\end{array}
\qquad
\begin{array}{c}
\text{REDPUSHARR} \\
\frac{c \text{ is fresh}}{G \langle \lambda(x : \tau) M \rangle N \rightsquigarrow_{\beta} (\pi(c)G) \langle M[x/c \langle N \rangle] \rangle}
\end{array}$$

However, pursuing this path while preserving confluence, typing, subject reduction, and progress is one of the main difficulties of coercion decomposition. The other main difficulty is strong normalization. We cannot reify coercion projections as lambda terms, so our current proof of termination does not work. We may argue that strong normalization is not necessary, but it is only true for β -reduction, because the ι -normalization is used in the proof of soundness. Said otherwise, we need ι -normalization, but we do not need to prove the strong normalization of the implicit reduction.

In the next chapter, we extend System F_t^P with unrestricted coercion abstraction. We only define the implicit version of this extension, called System F_{cc} . This avoids defining the explicit

reduction and as a consequence proving that the ι -reduction normalizes. More precisely, we do not need to exhibit a ι -reduction strategy to clean a redex in order to prove the backward simulation. We also avoid defining typing rules for coercion decomposition. There is however an underlying explicit version similar to System F_t^p , but missing backward simulation. We do not know if it is possible to define an explicit version with subject reduction (see Section 5.7).

Chapter 5

An implicit calculus of coercions: System F_{cc}

In this chapter we define a type system, called System F_{cc} , with unrestricted but safe erasable coercion abstractions and unrestricted but safe recursive types. Coercion abstraction is featured with coherent polymorphism which implies the consistency of abstract coercions. Like for **Constraint ML** but unlike for other type systems, we only give an implicit version for System F_{cc} , since an explicit version is hard to define and prove correct in the very general setting (see Section 4.6.2). The implicit version is proved sound with a step-index semantics. We developed a novel approach to step-index proofs, because we use a strong reduction setting while usual step-index proofs are done in a weak reduction setting. In particular, we put indices inside terms instead of aside them (see Section 5.2.1).

In this chapter, our approach to type systems is more semantic based. We define three notions: kinds which are interpreted as sets of mathematical objects, types which are interpreted as mathematical objects (the unit object, pairs of objects, or sets of terms for instance), and propositions interpreted as mathematical propositions.

We show that System F_{cc} satisfies the soundness property and is strongly normalizing in the absence of recursive types (Section 5.3). We show that System F_l^p can be seen as a sub-language of System F_{cc} (Section 5.4.2). And we describe how **Constraint ML** is included in System F_{cc} (Section 5.4.3). The soundness and normalization proofs are actually done in Coq, but with minor differences with the paper that are detailed in Section 5.6.

5.1 Definition

The definition of syntactic classes are given on Figure 5.1. We define a syntactical class, called kinds and written κ , to classify types. Kinds contain the star kind \star to classify sets of terms, the unit kind 1 to classify the unit mathematical object, product kinds $\kappa \times \kappa$ to classify mathematical pairs, and constrained kinds $\{\alpha : \kappa \mid P\}$ to constrain the set κ to its elements α satisfying the proposition P . For instance, the constrained kind $\{\alpha : \star \mid (\emptyset \vdash \alpha) \triangleright \tau\}$ corresponds to the set of types σ smaller than $\tau[\alpha/\sigma]$, which also corresponds to the domain of the upper bounded polymorphic type of System $F_{<}$: or F_l^p .

As usual, we define the syntactical class of types, written τ or σ . Types contain variables α , the unit object $\langle \rangle$, pairs $\langle \tau, \tau \rangle$, and projections **fst** τ and **snd** τ . Types additionally contain the arrow types $\tau \rightarrow \tau$, product types $\tau \times \tau$, polymorphic types $\forall(\alpha : \kappa) \tau$, recursive types $\mu\alpha \tau$,

$\kappa ::= \star \mid 1 \mid \kappa \times \kappa \mid \{\alpha : \kappa \mid P\}$	Kinds
$\tau, \sigma ::= \alpha \mid \langle \rangle \mid \langle \tau, \tau \rangle \mid \text{fst } \tau \mid \text{snd } \tau$	Types
$\quad \mid \tau \rightarrow \tau \mid \tau \times \tau \mid \forall(\alpha : \kappa) \tau \mid \mu\alpha \tau \mid \perp \mid \top$	
$P ::= \top \mid P \wedge P \mid (\Sigma \vdash \tau) \triangleright \tau \mid \exists \kappa \mid \forall(\alpha : \kappa) P$	Propositions
$\Sigma ::= \emptyset \mid \Sigma, (\alpha : \kappa)$	Type envs
$\Theta ::= \emptyset \mid \Theta, P$	Proof envs
$\Gamma ::= \emptyset \mid \Gamma, (x : \tau)$	Term envs
$\text{rec} ::= \text{NE} \mid \text{WF}$	well-foundedness signs

Figure 5.1: System F_{cc} syntax

the bottom type \perp , and the top type \top . The types of the previous sentence are interpreted as sets of terms. We define the projections of pairs equal to the associated components: $\text{fst } \langle \tau, \sigma \rangle$ is equal to τ and $\text{snd } \langle \tau, \sigma \rangle$ is equal to σ . Type equality is reflexive, symmetrical, transitive, and congruent for all syntactical constructs (see the Coq inductive `typesystem.jeq` for more details). We could also have added $\mu\alpha \tau = \tau[\alpha/\mu\alpha \tau]$ in the equality rules, but this rule needs typing conditions to ensure that τ is actually of kind \star and we do not have typing conditions in our syntactical equality rules. We could have defined recursive types as $\mu(\alpha : \kappa)\tau$ and then have the given equality, but the semantics of recursive types for arbitrary kinds is more involved. As a consequence, we have the recursive type equality at the coercion level, where type kinding is available. See Section 6.1.4 for discussion about arbitrary kind recursive types.

Instead of having only proofs of inclusion between sets of terms, that we call coercions, we define a general notion of propositions and define a syntactical class, written P , for them. Propositions contain the true proposition \top , conjunctions $P \wedge P$, coercions $(\Sigma \vdash \tau) \triangleright \tau$, coherence propositions $\exists \kappa$, and polymorphic propositions $\forall(\alpha : \kappa) P$. These propositions are what we can usually find in type systems based on constraints, which gives an intuition of how System F_{cc} compares to Constraint ML. Notice that the coherence proposition of the constrained kind $\{\alpha : \kappa \mid P\}$, namely $\exists\{\alpha : \kappa \mid P\}$, gives the usual existential proposition. Because coherence corresponds to habitation and a type τ is in the constrained kind $\{\alpha : \kappa \mid P\}$ if it is in κ and satisfies P by definition.

We define environments for types, proofs, and terms. Type environments, written Σ , are lists of bindings of the form $(\alpha : \kappa)$ which binds the type variable α to its kind κ . Notice that since kinds depend upon types, they may refer to type variables, and thus type environments are dependent and the order really matters. Proof environments, written Θ , collect which propositions hypotheses are accessible and hence they are lists of propositions. We don't need a variable to designate the proof since we only define the implicit version of our type system. Finally, term environments, written Γ , are lists of bindings of the form $(x : \tau)$ binding the term variable x to its type τ . The order in proof and term environments does not matter, but we still keep them as lists since this is how they are represented in the Coq proof.

Finally, we define the well-foundedness signs NE and WF , which we designate with the meta-variable rec . These signs are used to know which type functors are non-expansive, for sign NE , or well-founded, for sign WF .

Similarly to System F_{rec} , we define a well-foundedness judgment, written $\alpha \mapsto \tau : \text{rec}$, which is interpreted as the functor associating the type variable α to type τ is non-expansive (resp. well-founded) if rec is NE (resp. WF). The rules of this judgment are given on Figure 5.2.

$$\begin{array}{c}
\text{RECVAR} \\
\alpha \mapsto \alpha : \text{NE}
\end{array}
\qquad
\frac{\text{REARR} \quad \alpha \mapsto \tau : \text{NE} \quad \alpha \mapsto \sigma : \text{NE}}{\alpha \mapsto \tau \rightarrow \sigma : \text{WF}}
\qquad
\frac{\text{REPROD} \quad \alpha \mapsto \tau : \text{NE} \quad \alpha \mapsto \sigma : \text{NE}}{\alpha \mapsto \tau \times \sigma : \text{WF}}$$

$$\frac{\text{REFOR} \quad \alpha \mapsto \tau : \text{rec} \quad \alpha \notin \text{fv}(\kappa)}{\alpha \mapsto \forall(\beta : \kappa) \tau : \text{rec}}
\qquad
\frac{\text{RECMU} \quad \alpha \mapsto \tau : \text{rec} \quad \beta \mapsto \tau : \text{WF}}{\alpha \mapsto \mu\beta \tau : \text{rec}}$$

$$\frac{\text{RECWF} \quad \alpha \notin \text{fv}(\tau)}{\alpha \mapsto \tau : \text{WF}}
\qquad
\frac{\text{RECNE} \quad \alpha \mapsto \tau : \text{WF}}{\alpha \mapsto \tau : \text{NE}}$$

Figure 5.2: System F_{cc} well-foundedness judgment relation

The identity functor $\alpha \mapsto \alpha$ is non-expansive according to rule `RECVAR`. The arrow and product types are well-founded with respect to the type variable α , if their components are non-expansive with respect to α , according to rules `REARR` and `REPROD`. Rule `REFOR` tells that the polymorphic type $\forall(\beta : \kappa) \tau$ has the same well-foundedness sign than τ with respect to the type variable α which has to be different from β . Besides, the kind κ must not mention the type variable α . Similarly, rule `RECMU` tells that the recursive type $\mu\beta \tau$ preserves the well-foundedness sign of its body τ with respect to type variable α . The additional condition is that $\mu\beta \tau$ has to be a well-formed recursive type, in other words, the type τ has to be well-founded with respect to the type variable β . The last two rules are generic ways to prove well-foundedness or non-expansiveness. Constant functors are well-founded by rule `RECWF` and well-founded functors are non-expansive by rule `RECNE`.

Because most of the judgments in System F_{cc} are mutually recursive, we give a snapshot of all possible judgments with their intuition and mathematical interpretation. The only particularities are the well-foundedness judgment, which we just defined, because it has no dependencies, and the term judgment, because no judgment depends on it. We write $\Sigma \Vdash \kappa$ when the kind κ is well-formed under the type environment Σ . Similarly, we write $\Sigma \Vdash P$, $\Sigma \Vdash \Sigma$, and $\Sigma \Vdash \Theta$, for proposition, type environment, and proposition environment well-formedness, respectively. Well-formedness judgments have no mathematical interpretation. They are only used for extraction (Lemma 95), which we describe when presenting each judgment.

The kind judgment $\Sigma \vdash \kappa$ means that the kind κ is well-formed and coherent under Σ . It is interpreted as the mathematical proposition that the interpretation of κ under interpretations of Σ is inhabited. The type judgment $\Sigma \vdash \tau : \kappa$ means that the type τ has kind κ under environment Σ . Its interpretation is that the interpretation of τ is an element of the interpretation of κ under interpretations of Σ . When Σ is well-formed, we can extract from $\Sigma \vdash \tau : \kappa$ that κ is well-formed under Σ . The proposition judgment $\Sigma; \Theta_0; \Theta_1 \vdash P$ means that the proposition P holds under Σ with accessible coinduction hypotheses Θ_0 and protected coinduction hypotheses Θ_1 . We use coinduction for recursive types (see Section 5.4.4). Our notion of productivity comes from computational types (see rule `COERARR` for example). The mathematical interpretation of the proposition judgment is involved and uses step-indices (see Lemma 100). When Σ is well-formed and both Θ_0 and Θ_1 are well-formed under Σ , we can extract from $\Sigma; \Theta_0; \Theta_1 \vdash P$ that P is well-formed under Σ .

$$\frac{\text{KIND} \quad \Sigma; \emptyset; \emptyset \vdash \exists \kappa \quad \Sigma \Vdash \kappa}{\Sigma \vdash \kappa}$$

Figure 5.3: System F_{cc} kind judgment relation

$$\begin{array}{c} \text{TYPECONV} \\ \frac{\Sigma \vdash \tau : \kappa \quad \kappa = \kappa' \quad \Sigma \Vdash \kappa'}{\Sigma \vdash \tau : \kappa'} \\ \\ \text{TYPEVAR} \\ \frac{(\alpha : \kappa) \in \Sigma}{\Sigma \vdash \alpha : \kappa} \\ \\ \text{TYPEARR} \\ \frac{\Sigma \vdash \tau : \star \quad \Sigma \vdash \sigma : \star}{\Sigma \vdash \tau \rightarrow \sigma : \star} \\ \\ \text{TYPEPROD} \\ \frac{\Sigma \vdash \tau : \star \quad \Sigma \vdash \sigma : \star}{\Sigma \vdash \tau \times \sigma : \star} \\ \\ \text{TYPEFOR} \\ \frac{\Sigma \Vdash \kappa \quad \Sigma, (\alpha : \kappa) \vdash \tau : \star}{\Sigma \vdash \forall (\alpha : \kappa) \tau : \star} \\ \\ \text{TYPEMU} \\ \frac{\alpha \mapsto \tau : \text{WF} \quad \Sigma, (\alpha : \star) \vdash \tau : \star}{\Sigma \vdash \mu \alpha \tau : \star} \\ \\ \text{TYPEBOT} \quad \text{TOTYPETOP} \quad \text{TYPEUNIT} \\ \Sigma \vdash \perp : \star \quad \Sigma \vdash \top : \star \quad \Sigma \vdash \langle \rangle : 1 \\ \\ \text{TYPEPAIR} \\ \frac{\Sigma \vdash \tau_1 : \kappa_1 \quad \Sigma \vdash \tau_2 : \kappa_2}{\Sigma \vdash \langle \tau_1, \tau_2 \rangle : \kappa_1 \times \kappa_2} \\ \\ \text{TYPEFST} \\ \frac{\Sigma \vdash \tau : \kappa_1 \times \kappa_2}{\Sigma \vdash \text{fst } \tau : \kappa_1} \\ \\ \text{TYPESND} \\ \frac{\Sigma \vdash \tau : \kappa_1 \times \kappa_2}{\Sigma \vdash \text{snd } \tau : \kappa_2} \\ \\ \text{TYPEPACK} \\ \frac{\Sigma, (\alpha : \kappa) \Vdash P \quad \Sigma \vdash \tau : \kappa \quad \Sigma; \emptyset; \emptyset \vdash P[\alpha/\tau]}{\Sigma \vdash \tau : \{\alpha : \kappa \mid P\}} \\ \\ \text{TYPEUNPACK} \\ \frac{\Sigma \vdash \tau : \{\alpha : \kappa \mid P\}}{\Sigma \vdash \tau : \kappa} \end{array}$$

Figure 5.4: System F_{cc} type judgment relation

We define in Figure 5.3 when a kind is coherent. A kind κ is coherent with respect to a type environment Σ , written $\Sigma \vdash \kappa$, when the kind coherence proposition $\exists \kappa$ is satisfied under Σ . Rule **KIND** is the only rule for the kind coherence judgment. This means that the kind judgment is just a notation for the kind coherence proposition (see rule **PROPEXI** in Figure 5.5). Notice that we give no coinduction hypotheses to the proposition, which explains the two empty proposition environments in the premise.

A type τ has kind κ under type environment Σ , written $\Sigma \vdash \tau : \kappa$, if it satisfies the rules given in Figure 5.4. Rule **TYPECONV** permits to convert from κ to κ' the kind of a type τ , given both kinds are equal and the final kind κ' is well-formed under the given environment.

Rule **TYPEVAR** tells that the type variable α has kind κ under the type environment Σ , if α is bound to κ in Σ . Arrow and product types have the star kind if their components do too, by rules **TYPEARR** and **TYPEPROD** respectively. The type environment remains the same since none of these two type constructs binds type variables. One example for type binding is rule **TYPEFOR** for polymorphic types. The polymorphic type $\forall (\alpha : \kappa) \tau$ has the star kind under the type environment Σ if the body type τ has the star kind under the environment Σ extended with the type binding $(\alpha : \kappa)$ associating the type variable α to its well-formed kind κ . We do not ask κ to be coherent, because this condition is only necessary for the type generalization coercion rule **COERGEN** in Figure 5.6.

Rule **TYPEMU** tells that the recursive type $\mu \alpha \tau$ has the star kind if its body has the star kind under the same environment extended with α bound to the star kind. First, notice that recursive types are only of the star kind (relaxing this restriction is discussed in Section 6.1.4).

$$\begin{array}{c}
\text{PROP CONV} \\
\frac{\Sigma; \Theta_0; \Theta_1 \vdash P \quad P = P' \quad \Sigma \Vdash P'}{\Sigma; \Theta_0; \Theta_1 \vdash P'} \\
\\
\text{PROP VAR} \\
\frac{P \in \Theta_0}{\Sigma; \Theta_0; \Theta_1 \vdash P} \\
\\
\text{PROP TRUE} \\
\Sigma; \Theta_0; \Theta_1 \vdash \top \\
\\
\text{PROP PAIR} \\
\frac{\Sigma; \Theta_0; \Theta_1 \vdash P_1 \quad \Sigma; \Theta_0; \Theta_1 \vdash P_2}{\Sigma; \Theta_0; \Theta_1 \vdash P_1 \wedge P_2} \\
\\
\text{PROP FST} \\
\frac{\Sigma; \Theta_0; \Theta_1 \vdash P_1 \wedge P_2}{\Sigma; \Theta_0; \Theta_1 \vdash P_1} \\
\\
\text{PROP SND} \\
\frac{\Sigma; \Theta_0; \Theta_1 \vdash P_1 \wedge P_2}{\Sigma; \Theta_0; \Theta_1 \vdash P_2} \\
\\
\text{PROP EXI} \\
\frac{\Sigma \vdash \tau : \kappa}{\Sigma; \Theta_0; \Theta_1 \vdash \exists \kappa} \\
\\
\text{PROP GEN} \\
\frac{\Sigma \Vdash \kappa \quad \Sigma, (\alpha : \kappa); \Theta_0; \Theta_1 \vdash P \quad \alpha \notin \text{fv}(\Theta_0, \Theta_1)}{\Sigma; \Theta_0; \Theta_1 \vdash \forall (\alpha : \kappa) P} \\
\\
\text{PROP INST} \\
\frac{\Sigma; \Theta_0; \Theta_1 \vdash \forall (\alpha : \kappa) P \quad \Sigma \vdash \tau : \kappa}{\Sigma; \Theta_0; \Theta_1 \vdash P[\alpha/\tau]} \\
\\
\text{PROP RES} \\
\frac{\Sigma \vdash \tau : \{\alpha : \kappa \mid P\}}{\Sigma; \Theta_0; \Theta_1 \vdash P[\alpha/\tau]} \\
\\
\text{PROP FIX} \\
\frac{\Sigma \Vdash P \quad \Sigma; \Theta_0; \Theta_1, P \vdash P}{\Sigma; \Theta_0; \Theta_1 \vdash P} \\
\\
\text{PROP COER} \\
\frac{\Sigma; \Theta_0; \Theta_1 \vdash^* (\Sigma' \vdash \tau') \triangleright \tau \quad \Sigma, \Sigma' \vdash \tau' : \star}{\Sigma; \Theta_0; \Theta_1 \vdash (\Sigma' \vdash \tau') \triangleright \tau}
\end{array}$$

Figure 5.5: System F_{cc} proposition judgment relation

Then, a recursive type has also to be well-founded as a functor to be well-kinded, which is ensured by hypothesis $\alpha \mapsto \tau : \text{WF}$. The bottom and top types have kind star according to rules TYPEBOT and TYPETOP .

By rule TYPEUNIT , the unit type has the unit kind. The pair of types τ_1 and τ_2 has the product kind $\kappa_1 \times \kappa_2$ under type environment Σ , given type τ_1 has kind κ_1 and type τ_2 has kind κ_2 under Σ . This rule is similar to the rule for pairs in terms. And like for terms, type projections extract their kind from their premise. The first projection $\text{fst } \tau$ and second projection $\text{snd } \tau$ have kinds κ_1 and κ_2 respectively, if the type τ has the product kind $\kappa_1 \times \kappa_2$ under the same environment.

Finally, the last two rules are particular since they just change the kind of a type leaving the type unchanged. In this sense they are similar to coercions but at the kind level (see Section 6.1.8). This is actually called subkinding since in our case the environment are not modified. Rule TYPEPACK tells that if a type τ has kind κ and satisfies the well-formed property P , then it also has the constrained kind $\{\alpha : \kappa \mid P\}$ of types of kind κ satisfying P . Rule TYPEUNPACK acts the opposite way: if type τ has constrained kind $\{\alpha : \kappa \mid P\}$, then it also has kind κ by forgetting the fact that it also satisfies the property P . The constrained kind can be seen as a dependent sum of a kind and a proposition. It is built with a type and a proof that this type satisfies a proposition by rule TYPEPACK . It can be eliminated on its first component by rule TYPEUNPACK . And it can also be eliminated on its second component by rule PROPRES in Figure 5.5.

The proposition judgment contains logical rules given in Figure 5.5. It is written $\Sigma; \Theta_0; \Theta_1 \vdash P$. It tells that the proposition P is true under the type environment Σ , given additional coinduction hypotheses in Θ_0 and Θ_1 . The hypotheses in Θ_0 are accessible while those in Θ_1 will be accessible only after some productive proof step has been done. Productive proof steps are

exclusively done in the η -expansion coercion rules of computational types (see Figure 5.6). Coinduction hypotheses are extended in the coinduction proof step. These are the only rule altering the coinduction hypotheses.

Similar to rule `TYPECONV` for types, we can convert the proof of a proposition from P to P' under the same environment, given that P is equal to P' and P' is well-formed under the same environment. This rule is named `PROP_CONV` and given in Figure 5.5.

Rule `PROP_VAR` tells that a proposition P is true if it is present in the accessible coinduction hypotheses Θ_0 . The true proposition is always true, according to rule `PROP_TRUE`. Rule `PROP_PAIR` tells that the conjunction proposition $P_1 \wedge P_2$ is true if its components P_1 and P_2 are true under the same environments. If a conjunction proposition $P_1 \wedge P_2$ is true, then its components P_1 and P_2 are also true under the same environments by rules `PROP_FST` and `PROP_SND` respectively. These conjunction rules resemble those of the product type and the product kind, except that we are not interested in their proof terms. A coherence proposition $\exists \kappa$ is true if its kind κ is inhabited, which is our notion of coherence as explained above.

Rule `PROP_GEN` and `PROP_INST` deal with proposition polymorphism. The polymorphic proposition $\forall(\alpha : \kappa) P$ holds under the type environment Σ and the coinduction hypotheses Θ_0 and Θ_1 , if its inner proposition P holds under the extended type environment $\Sigma, (\alpha : \kappa)$, given that κ is well-formed under Σ and α is not free in the coinduction hypotheses. To eliminate such polymorphic proposition, we supply it with a type of the given kind. If $\forall(\alpha : \kappa) P$ holds and τ has kind κ , then the substituted proposition $P[\alpha/\tau]$ holds under the same environment.

Rule `PROP_RES` allows to extract from a type of a constrained kind the fact that the property holds for this type. This is the second projection rule of the constrained kind—the first being rule `TYPE_UNPACK`. If τ has kind $\{\alpha : \kappa \mid P\}$, then $P[\alpha/\tau]$ is true under the same type environment and any coinduction environments. A rule we do not usually find in type systems is rule `PROP_FIX`: it allows to prove a proposition by coinduction. If P is true assuming P in the protected coinduction environment, then P is true without this additional hypothesis. We can derive from this general rule the usual rules about recursive types we can find in other type systems (see Section 5.4.4). Finally, rule `PROP_COER` permits to see a coercion proposition as a coercion, given that the inner type has kind star.

The coercion judgment is written $\Sigma; \Theta_0; \Theta_1 \vdash^* (\Sigma' \vdash \tau') \triangleright \tau$. Notice how the turn-style is annotated with a star in order to differentiate the coercion judgment from the proposition judgment of the coercion proposition. This difference matters only for the extraction lemma: these two judgments have the same interpretation. Coercion rules are given in Figure 5.6. They resemble very closely to what we already have seen in System F_t^p . A coercion proposition $(\Sigma' \vdash \tau') \triangleright \tau$ is well-formed under Σ , if its binding environment Σ' is coherent under Σ (all its kinds are coherent), its argument type τ' has kind star under the extended environment Σ, Σ' , and its return type τ has kind star under Σ . See rule `WFPCOER` in Figure 5.8 for the formal statement.

Coercions can be coercion propositions by rule `COER_PROP`. Coercions are closed by reflexivity and transitivity since they are interpreted as inclusions. Rule `COER_REFL` in Figure 5.6 tells that a type τ is smaller than itself. Rule `COER_TRANS` tells that τ_1 extended with Σ_2, Σ_1 is smaller than τ_3 under Σ , if τ_2 extended with Σ_2 is smaller than τ_3 under Σ and τ_1 extended with Σ_1 is smaller than τ_2 under Σ, Σ_2 . The coinduction environments remain the same, but the type variables bound in Σ_2 should not be free in the coinduction environments. If τ extended with Σ' is smaller than σ and τ does not mention the type variables in Σ' , then τ is smaller than σ by rule `COER_WEAK`.

Rules `COER_ARR` and `COER_PROD` are about η -expansions of computational types. Compu-

$$\begin{array}{c}
\text{COERPROP} \\
\frac{\Sigma; \Theta_0; \Theta_1 \vdash (\Sigma \vdash \tau') \triangleright \tau}{\Sigma; \Theta_0; \Theta_1 \vdash^* (\Sigma \vdash \tau') \triangleright \tau}
\end{array}
\qquad
\begin{array}{c}
\text{COERREFL} \\
\Sigma; \Theta_0; \Theta_1 \vdash^* (\emptyset \vdash \tau) \triangleright \tau
\end{array}$$

$$\begin{array}{c}
\text{COERTRANS} \\
\frac{\Sigma, \Sigma_2; \Theta_0; \Theta_1 \vdash^* (\Sigma_1 \vdash \tau_1) \triangleright \tau_2 \quad \Sigma; \Theta_0; \Theta_1 \vdash^* (\Sigma_2 \vdash \tau_2) \triangleright \tau_3}{\Sigma; \Theta_0; \Theta_1 \vdash^* (\Sigma_2, \Sigma_1 \vdash \tau_1) \triangleright \tau_3}
\end{array}
\qquad
\begin{array}{c}
\text{COERWEAK} \\
\frac{\text{dom}(\Sigma') \cap \text{fv}(\tau) = \emptyset \quad \Sigma; \Theta_0; \Theta_1 \vdash^* (\Sigma' \vdash \tau) \triangleright \sigma}{\Sigma; \Theta_0; \Theta_1 \vdash^* (\emptyset \vdash \tau) \triangleright \sigma}
\end{array}$$

$$\begin{array}{c}
\text{COERARR} \\
\frac{\text{dom}(\Sigma') \cap \text{fv}(\Theta_0, \Theta_1, \tau) = \emptyset \quad \Sigma \vdash \tau : \star \quad \Sigma, \Sigma'; \Theta_0, \Theta_1; \emptyset \vdash^* (\emptyset \vdash \tau) \triangleright \tau' \quad \Sigma; \Theta_0, \Theta_1; \emptyset \vdash^* (\Sigma' \vdash \sigma') \triangleright \sigma}{\Sigma; \Theta_0; \Theta_1 \vdash^* (\Sigma' \vdash \tau' \rightarrow \sigma') \triangleright \tau \rightarrow \sigma}
\end{array}$$

$$\begin{array}{c}
\text{COERPROD} \\
\frac{\Sigma; \Theta_0, \Theta_1; \emptyset \vdash^* (\Sigma' \vdash \tau') \triangleright \tau \quad \Sigma; \Theta_0, \Theta_1; \emptyset \vdash^* (\Sigma' \vdash \sigma') \triangleright \sigma}{\Sigma; \Theta_0; \Theta_1 \vdash^* (\Sigma' \vdash \tau' \times \sigma') \triangleright \tau \times \sigma}
\end{array}$$

$$\begin{array}{c}
\text{COERGEN} \\
\frac{\Sigma \vdash \kappa}{\Sigma; \Theta_0; \Theta_1 \vdash^* (\emptyset, (\alpha : \kappa) \vdash \tau) \triangleright \forall(\alpha : \kappa) \tau}
\end{array}
\qquad
\begin{array}{c}
\text{COERINST} \\
\frac{\Sigma \vdash \sigma : \kappa}{\Sigma; \Theta_0; \Theta_1 \vdash^* (\emptyset \vdash \forall(\alpha : \kappa) \tau) \triangleright \tau[\alpha/\sigma]}
\end{array}$$

$$\begin{array}{c}
\text{COERUNFOLD} \\
\Sigma; \Theta_0; \Theta_1 \vdash^* (\emptyset \vdash \mu\alpha \tau) \triangleright \tau[\alpha/\mu\alpha \tau]
\end{array}
\qquad
\begin{array}{c}
\text{COERFOLD} \\
\frac{\Sigma \vdash \mu\alpha \tau : \star}{\Sigma; \Theta_0; \Theta_1 \vdash^* (\emptyset \vdash \tau[\alpha/\mu\alpha \tau]) \triangleright \mu\alpha \tau}
\end{array}$$

$$\begin{array}{c}
\text{COERTOP} \\
\Sigma; \Theta_0; \Theta_1 \vdash^* (\emptyset \vdash \tau) \triangleright \top
\end{array}
\qquad
\begin{array}{c}
\text{COERBOT} \\
\frac{\Sigma \vdash \tau : \star}{\Sigma; \Theta_0; \Theta_1 \vdash^* (\emptyset \vdash \perp) \triangleright \tau}
\end{array}$$

Figure 5.6: System F_{cc} coercion judgment relation

tational types are the only types that need η -expansion coercion rules, because η -expansion coercion rules for erasable types are derivable from their introduction and elimination rules. Notice however that the η -expansion of recursive types uses coinduction (see Section 5.4.4). As η -expansion rules of computational types, rules COERARR and COERPROD are productive proofs for the coinduction. This is why the protected coinduction environment, namely Θ_1 , of the conclusion is accessible in the premises. If σ' extended with Σ' is smaller than σ under Σ and τ , which has kind star under Σ , is smaller than τ' under Σ, Σ' , then $\tau' \rightarrow \sigma'$ extended with Σ' is smaller than $\tau \rightarrow \sigma$ by rule COERARR. Similarly if τ' extended with Σ' is smaller than τ and σ' extended with Σ' is smaller than σ , then $\tau' \times \sigma'$ extended with Σ' is smaller than $\tau \times \sigma$.

Rules COERGEN and COERINST are about coherent polymorphism. The type τ extended with the type binding $(\alpha : \kappa)$ is smaller than the polymorphic type $\forall(\alpha : \kappa) \tau$, given that κ is coherent. This coherence condition is an important feature of System F_{cc} : polymorphism is erasable (and hence a coercion proposition) only if it is coherent. Notice that the usual polymorphism for kind star is always coherent, because the star kind is trivially inhabited

$$\begin{array}{c}
\text{TENVNIL} \\
\Sigma \vdash \emptyset
\end{array}
\qquad
\frac{\text{TENVCONS} \quad \Sigma \vdash \Sigma' \quad \Sigma, \Sigma' \vdash \kappa}{\Sigma \vdash \Sigma', (\alpha : \kappa)}$$

Figure 5.7: System F_{cc} type environment judgment relation

$$\begin{array}{c}
\text{WFKSTAR} \quad \text{WFKONE} \quad \frac{\text{WFKPROD} \quad \Sigma \Vdash \kappa_1 \quad \Sigma \Vdash \kappa_2}{\Sigma \Vdash \kappa_1 \times \kappa_2} \quad \frac{\text{WFKRES} \quad \Sigma \Vdash \kappa \quad \Sigma, (\alpha : \kappa) \Vdash P}{\Sigma \Vdash \{\alpha : \kappa \mid P\}} \quad \text{WFPTRUE} \\
\Sigma \Vdash \star \quad \Sigma \Vdash 1
\end{array}$$

$$\frac{\text{WFPAND} \quad \Sigma \Vdash P_1 \quad \Sigma \Vdash P_2}{\Sigma \Vdash P_1 \wedge P_2} \quad \frac{\text{WFPcoER} \quad \Sigma \vdash \Sigma' \quad \Sigma, \Sigma' \vdash \tau' : \star \quad \Sigma \vdash \tau : \star}{\Sigma \Vdash (\Sigma' \vdash \tau') \triangleright \tau} \quad \frac{\text{WFPExI} \quad \Sigma \Vdash \kappa}{\Sigma \Vdash \exists \kappa} \quad \frac{\text{WFPFOR} \quad \Sigma, (\alpha : \kappa) \Vdash P}{\Sigma \Vdash \forall (\alpha : \kappa) P}$$

$$\frac{\text{WFHNIL} \quad \Sigma \Vdash \emptyset}{\Sigma \Vdash \emptyset} \quad \frac{\text{WFHCONS} \quad \Sigma \Vdash \Sigma' \quad \Sigma, \Sigma' \Vdash \kappa}{\Sigma \Vdash \Sigma', (\alpha : \kappa)} \quad \frac{\text{WFYNIL} \quad \Sigma \Vdash \emptyset}{\Sigma \Vdash \emptyset} \quad \frac{\text{WFYCONS} \quad \Sigma \Vdash \Theta \quad \Sigma \Vdash P}{\Sigma \Vdash \Theta, P}$$

Figure 5.8: System F_{cc} well-formedness relation

(with the top or bottom type for instance). Notice also that bounded polymorphism is also trivially coherent (see Section 5.4.2). Without this coherence condition we could give the term $27 + \text{true}$ (the addition of the integer constant 27 and the boolean constant `true`) the type $\forall(\alpha : \{\alpha : 1 \mid (\emptyset \vdash \text{Bool}) \triangleright \text{Int}\}) \text{Int}$. We simply use rule `PROPRES` to extract from the abstract type variable α the fact that booleans can be coerced to integers. This kind of feature is actually useful to deal with GADTs and we present it as a Coq-proved extension in Section 5.5. Similarly to type generalization, type instantiation is a coercion and the polymorphic type $\forall(\alpha : \kappa) \tau$ is smaller than its instantiation $\tau[\alpha/\sigma]$ with type σ of kind κ .

Rules `COERUNFOLD` and `COERFOLD` are about recursive types. The well-kinded recursive type $\mu\alpha \tau$ is equivalent (*i.e.* smaller and bigger than) to its unfolding $\tau[\alpha/\mu\alpha \tau]$. Finally, rules `COERBOT` and `COERTOP` are about extrema. The bottom type \perp is smaller than all types of kind star and the top type \top is bigger than all types of kind star.

The type environment Σ' is coherent under Σ , written $\Sigma \vdash \Sigma'$, if its type variables are bound at most once and the associated kinds are coherent in their preceding type environment. The rules for type environment coherence are given in Figure 5.7. The empty type environment \emptyset is always coherent according to rule `TENVNIL`. The extended environment $\Sigma', (\alpha : \kappa)$ is coherent under Σ if the type variable α is not already bound in Σ, Σ' and the kind κ is coherent under the extended environment Σ, Σ' .

We use the notation $\Sigma \Vdash \cdot$ for the well-formedness judgment of all syntactical classes subject to well-formedness rules: kinds, propositions, type environments, and coinduction environments. The well-formedness judgment has no mathematical interpretation. It is only use for the extraction lemma (Lemma 95). The well-formedness rules are given in Figure 5.8. An object is well-formed if its components are well-formed in an extended type environment when the object behaves as a binder. The only particular rule is for the coercion proposition, namely rule `WFPCOER`, because it asks for coherence and well-kindedness instead of well-

$$\begin{array}{c}
\text{ENVNIL} \\
\Sigma \vdash \emptyset
\end{array}
\qquad
\frac{\text{ENVCONS} \quad \Sigma \vdash \Gamma \quad \Sigma \vdash \tau : \star \quad x \notin \text{dom}(\Gamma)}{\Sigma \vdash \Gamma, (x : \tau)}$$

Figure 5.9: System F_{cc} term environment judgment relation

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

$$\begin{array}{c}
\frac{\text{TERMPAIR} \quad a : \Sigma; \Gamma \vdash \tau \quad b : \Sigma; \Gamma \vdash \sigma}{\langle a, b \rangle : \Sigma; \Gamma \vdash \tau \times \sigma}
\qquad
\frac{\text{TERMFST} \quad a : \Sigma; \Gamma \vdash \tau \times \sigma}{\text{fst } a : \Sigma; \Gamma \vdash \tau}
\qquad
\frac{\text{TERMSND} \quad a : \Sigma; \Gamma \vdash \tau \times \sigma}{\text{snd } a : \Sigma; \Gamma \vdash \sigma}$$

$$\frac{\text{TERMCOR} \quad \text{dom}(\Sigma') \cap \text{fv}(\Gamma) = \emptyset \quad a : \Sigma, \Sigma'; \Gamma \vdash \tau' \quad \Sigma; \emptyset; \emptyset \vdash^* (\Sigma' \vdash \tau') \triangleright \tau}{a : \Sigma; \Gamma \vdash \tau}$$

Figure 5.10: System F_{cc} term judgment relation

formedness. A coercion proposition $(\Sigma' \vdash \tau') \triangleright \tau$ is well-formed if Σ' is coherent and both τ and τ' have kind star under their respective environments.

A term environment Γ is valid under Σ , written $\Sigma \vdash \Gamma$, if all its term variables are bound at most once and their associated types have kind star under Σ . The rules are given in Figure 5.9. The empty term environment is valid under any type environment by rule ENVNIL. The extended term environment $\Gamma, (x : \tau)$ is valid if τ has kind star and the term variable x is not already bound in Γ , which has to be valid as well, according to rule ENVCONS.

Since we separated type environments and term environments, the term judgment changes a little. We write $a : \Sigma; \Gamma \vdash \tau$ when the term a has type τ under the type environment Σ and the term environment Γ . The rules for the term judgment are given on Figure 5.10. These are the usual STLC rules plus the additional coercion typing rule. The mathematical interpretation of this judgment is that for all Σ -instantiation and all Γ -substitution in this instantiation, the term a after substitution is in the type τ after instantiation. Syntactically, we can extract from the term judgment that the type has kind star, if the type environment is well-formed and the term environment is valid.

Rule TERMVAR tells that the term variable x has type τ under type environment Σ and term environment Γ if x is bound to τ in Γ . The term abstraction $\lambda x a$ has type $\tau \rightarrow \sigma$ under Σ and Γ if its body a has type σ under the extended term environment $\Gamma, (x : \tau)$ and the same type environment Σ by Rule TERMLAM. If, under type environment Σ and term environment Γ , the term a has type $\tau \rightarrow \sigma$ and b has type τ , then their application ab has type σ .

Rule TERMPAIR tells that if a has type τ and b has type σ under type environment Σ and term environment Γ , then the pair $\langle a, b \rangle$ has type $\tau \times \sigma$. The other way around, if a has type $\tau \times \sigma$, then its first projection $\text{fst } a$ has type τ and its second projection $\text{snd } a$ has type σ under the same environments.

Finally, the coercion rule `TERMCOER` changes the typing of term a from $\Sigma, \Sigma'; \Gamma \vdash \tau'$ to $\Sigma; \Gamma \vdash \tau$, given there is a coercion proposition $(\Sigma' \vdash \tau') \triangleright \tau$ under the type environment Σ . To respect scoping, the type environment Σ' should not bind type variables that are free in Γ . If this is the case, it suffices to rename these variables in Σ' and τ' with fresh ones.

Strengthening The strengthening rule is not admissible when we allow incoherent type environments, but it is still admissible with coherent type environments. It is sound to add the strengthening rule as follows:

$$\frac{\text{TYPESTR} \quad \Sigma \vdash \kappa' \quad \Sigma, (\alpha' : \kappa') \vdash \tau : \kappa \quad \alpha' \notin \text{fv}(\tau, \kappa)}{\Sigma \vdash \tau : \kappa}$$

$$\frac{\text{PROPSTR} \quad \Sigma \vdash \kappa \quad \Sigma, (\alpha : \kappa); \Theta_0; \Theta_1 \vdash P \quad \alpha \notin \text{fv}(\Theta_0, \Theta_1, P)}{\Sigma; \Theta_0; \Theta_1 \vdash P}$$

Rule `TYPESTR` is the elimination rule of the coherence proposition (see rule `PROPEXI`). Coherence proposition actually behaves like an existential: $\exists \kappa$ means that there is a type τ of kind κ . Since the kind judgment is actually a notation for the coherence proposition, the premise $\Sigma \vdash \kappa'$ is actually $\Sigma; \emptyset; \emptyset \vdash \exists \kappa'$ and we are eliminating this existential proposition. To do so, we use the second premise assuming a type α' of kind κ' to prove that τ has kind κ . Since neither τ or κ mention the witness type α' , we can conclude that τ has kind κ without the hypothesis about the coherence of κ' .

Again, similarly to the type judgment rule `TYPESTR`, we can eliminate the existential proposition when proving a proposition by rule `PROPSTR`, if the kind κ is coherent under Σ and the proposition P holds under the extended environment $\Sigma, (\alpha : \kappa)$ with coinduction hypotheses Θ_0 and Θ_1 given that the type variable α is not free in Θ_0 , Θ_1 , and P .

5.2 Semantics

A term is sound if none of its reductions lead to an error (see Figure 2.3). To avoid the negation, it is easier to reason with valid terms defined as the complement of Ω , *i.e.* terms that are not errors, which we write \mathcal{U} . To remember these notations, Ω is the omega greek letter usually used for errors, while \mathcal{U} looks like the letter v for valid terms. Notice also that they are 180 degrees rotation of one another because they are complement of one another. Hence, a term is sound if all its reduction paths lead to valid terms. Since this construction appears repeatedly, we define the *expansion* of a set of terms \mathbf{R} , which we write $(\rightsquigarrow^* \mathbf{R})$, the set of terms a such that any reduction path starting with a leads to a term in \mathbf{R} . The set \mathcal{S} of sound terms is the expansion $(\rightsquigarrow^* \mathcal{U})$ of valid terms.

Head normal forms Δ are terms whose root node is a constructor, *i.e.* abstractions $\lambda x a$ and pairs $\langle a, a \rangle$, while neutral terms ∇ are variables x , applications $a a$, and projections `fst` a and `snd` a . Notice that Δ and ∇ are complement of one another, *i.e.* terms are the disjoint union of Δ and ∇ . To remember which notation is associated to which notion, we can see Δ as stable (it has a large basis) and ∇ as unstable (it has a pointy basis). Stability, in our analogy, means that we have a head constructor which cannot change by reduction, while instability means that reduction can modify our root node, to finally become a head constructor. Notice

x, y	Term variables
$e, f ::= x^k \mid \lambda^k x e \mid (e e)^k \mid \langle e, e \rangle^k \mid \mathbf{fst}^k e \mid \mathbf{snd}^k e$	Terms

Figure 5.11: Indexed Calculus syntax

$E^k ::= \lambda^k x [] \mid ([] e)^k \mid (e [])^k \mid \langle [], e \rangle^k \mid \langle e, [] \rangle^k \mid \mathbf{fst}^k [] \mid \mathbf{snd}^k []$	Evaluation contexts
$s ::= E^k[s] \mid (\langle s, s \rangle^k s)^k \mid \mathbf{fst}^k (\lambda^k x s) \mid \mathbf{snd}^k (\lambda^k x s)$	Errors

Figure 5.12: Indexed Calculus notations

also that, as for errors and valid terms, head normal forms and neutrals are 180 degrees of one another because they are complement of one another.

Progress is a way to double-check the definition of the semantics, by defining values syntactically and checking that semantic values (irreducible valid terms) are syntactic values. Prevalues are simply neutral values.

Lemma 71 (Progress). *If $a \in \mathcal{U}$ and $a \not\rightsquigarrow$, then a is of the form v .*

Proof. By induction on a . For each case, either it is an error, or it can reduce, or it is a value. □

The converse is also true, *i.e.* values do not contain errors. However, this won't remain true when we restrict the strategy, *e.g.* to call-by-value weak reduction. In this case, redefining the grammar of values, progress still holds, but some grammatically well-formed values may contain “inaccessible” errors, such as errors occurring under an abstraction.

Type soundness states that well-typed terms are sound. We prove this by interpreting syntactic types as semantic types which are themselves sets of terms. However, since we allow general recursive types the evaluation of terms may not terminate. This is not a problem, since type soundness is not about termination, but ruling out unsound terms, which if they reach an error do so in a finite number of steps.

The idea of step-indexed techniques is to stop the reduction after a certain number of steps, as if some initially available fuel (the number of allowed reduction steps) had all been consumed. Since errors are necessarily reached after a finite number of steps, we may always detect errors with some finite but arbitrary large number of reduction steps.

However, there is a difficulty using step indices in a strong reduction setting, as the usual semantic of the arrow type is not stable by reduction. We solve this by including the fuel inside terms, called indexed terms, and block the reduction internally when terms do not have enough fuel, rather than control the number of reduction steps externally.

5.2.1 The Indexed Calculus

Terms of the Indexed Calculus are terms of the λ -calculus where each node is annotated with a natural number called its index (or fuel). They are written with letter e or f and formally defined on Figure 5.11. Indexed terms are variables x^k , abstractions $\lambda^k x e$, applications $(e e)^k$, pairs $\langle e, e \rangle^k$, and projections $\mathbf{fst}^k e$ and $\mathbf{snd}^k e$.

As for the λ -calculus, we define the indexed errors, written s , on Figure 5.12. They contain terms that contain an error under a context and terms that are the application of a pair or the projection of an abstraction, similarly to lambda (or undecorated) errors. The indices

$$\begin{array}{ll}
\lfloor x^k \rfloor_j = x^{kj} & \lfloor \langle e, f \rangle^k \rfloor_j = \langle \lfloor e \rfloor_j, \lfloor f \rfloor_j \rangle^{kj} \\
\lfloor \lambda^k x e \rfloor_j = \lambda^{kj} x \lfloor e \rfloor_j & \lfloor \mathbf{fst}^k e \rfloor_j = \mathbf{fst}^{kj} \lfloor e \rfloor_j \\
\lfloor (ef)^k \rfloor_j = (\lfloor e \rfloor_j \lfloor f \rfloor_j)^{kj} & \lfloor \mathbf{snd}^k e \rfloor_j = \mathbf{snd}^{kj} \lfloor e \rfloor_j
\end{array}$$

Figure 5.13: Lower function

are ignored in this definition. We write E^k for one-hole context with index k . We have all contexts since we are in strong reduction. Finally we define neutrals and head normal forms which are complement of one another according to the set of indexed terms. Neutrals are variables x^k , applications $(ee)^k$, and projections $\mathbf{fst}^k e$ and $\mathbf{snd}^k e$, while head normal forms are abstractions $\lambda^k x e$ and pairs $\langle e, e \rangle^k$.

We write ∇ for neutrals, Δ for head normal forms, Ω for indexed errors, and \mathfrak{U} for valid indexed terms, overloading the notation of the λ -calculus, which is never ambiguous because we use distinct meta-variables for lambda terms and indexed terms. These sets are closely related, indeed the set of neutral lambda terms and the set of neutral indexed terms are the same up to the erasure of indices, and similarly for head normal forms, errors, and valid terms.

Intuitively, indices indicate the maximum number of reduction steps allowed under the given node. We want this invariant to hold during reduction, so reduction has to update indices. Since redexes usually delete several nodes of the left-hand side and reorganize subterms for the right-hand side, we must remember the indices of the removed nodes one way or another on their subterms. The easiest solution is to take the minimum of indices of the redex nodes and lower their subterms with this minimum. To do so we use an auxiliary lowering function on indexed terms, written $\lfloor e \rfloor_j$, and defined on Figure 5.13 (Coq definition `Flanguage.lower`). We use concatenation of indices to denote the minimum of their values. This is not ambiguous since we never use multiplication of indices. Lowering simply changes all indices in the term e with their minimum with j . With this definition, the β -reduction rule now becomes: $((\lambda^{j+1} x e) f)^{k+1} \rightsquigarrow \lfloor e[x/f] \rfloor_{kj}$.

The capture avoiding substitution $e[x/f]$ of term f for variable x in the term e replaces in e all free occurrences x^j of x by $\lfloor f \rfloor_j$. The definition is generalized in the obvious way to simultaneous substitutions. We use letter γ to range over substitutions. The lowering of substituted occurrences is necessary to make substitution commute with the lowering function. In particular, renaming commutes with the lowering function.

Lemma 72. $\lfloor e[x/f] \rfloor_k = \lfloor e \rfloor_k[x/f] = \lfloor e \rfloor_k[x/\lfloor f \rfloor_k]$

Proof. The first equality is proved with the Coq lemma `lower_subst` and the last equality is proved with `subst_lower` in file `Flanguage.v`. \square

The reduction rules of the Indexed Calculus mimic those of the λ -calculus, but with some index manipulation, as described in Figure 5.14 (Coq inductive `Flanguage.red`). Reduction can only proceed when the indices on the nodes involved in the reduction are strictly positive; the indices are lowered after reduction by the minimum of the involved indices decremented by one. As a corollary, reduction cannot occur at or under a node with a null index. This applies both to head reduction rules (`REDAPP`, `REDFST`, and `REDSND`) and to reductions in an evaluation context (`RULE REDCTX`). That is, a head reduction can only be applied along a path of the form $E_1^{k_1}[\dots E_p^{k_p}[e]]$ when indices k_i 's are all strictly positive; they are all decremented after the reduction.

$$\begin{array}{c}
\text{REDCTX} \\
\frac{e \rightsquigarrow f}{E^{k+1}[e] \rightsquigarrow E^k[f]}
\end{array}
\qquad
\begin{array}{c}
\text{REDAPP} \\
((\lambda^{j+1} x e) f)^{k+1} \rightsquigarrow [e[x/f]]_{kj}
\end{array}$$

$$\begin{array}{c}
\text{REDFST} \\
\text{fst}^{k+1} \langle e, f \rangle^{j+1} \rightsquigarrow [e]_{kj}
\end{array}
\qquad
\begin{array}{c}
\text{REDSND} \\
\text{snd}^{k+1} \langle e, f \rangle^{j+1} \rightsquigarrow [f]_{kj}
\end{array}$$

Figure 5.14: Indexed Calculus reduction relation

For example, here is a decorated reduction of apply (the lambda term $\lambda x \lambda y x y$) applied to two terms e and f :

$$\rightsquigarrow \frac{(((\lambda^{k_3+1} x \lambda^{j_3} y (x^{j_3} y^{j_4})^{j_2}) e)^{k_2+1} f)^{k_1+1}}{((\lambda^{j_1 k_2 k_3} y ([e]_{j_3 k_2 k_3} y^{j_4 k_2 k_3})^{j_2 k_2 k_3}) f)^{k_1}}$$

Since the reduction happens under the external application, it must have some fuel $k_1 + 1$, which is decreased by one in the result. Then, for the redex to fire, the application must have some fuel $k_2 + 1$ as well as the abstraction $k_3 + 1$, which are both decreased by one and combined as $k_2 k_3$ to lower the result of the reduction. Before that, the term e , which has been substituted for x^{j_3} has been lowered to j_3 in the result. The important feature is that f has not been lowered, which is an important difference with what would happen with the traditional step-indexed approach when indices are outside terms.

Strong normalization By design, the Indexed Calculus is strongly normalizing, *i.e.* all reduction paths of all terms are finite. In particular, they are bounded by the index of their root node.

Lemma 73 (Strong normalization). *The indexed calculus is strongly normalizing.*

Proof. The proof is done in Coq lemma `Fnormalization.wf_der` and uses the measure that associates for each indexed term its root index. \square

5.2.2 Bisimulation

To show that reduction between undecorated terms and decorated terms coincide, we define $[e]$ (without subindex) the erasure of an indexed term e obtained by dropping all indices. We lift this function to sets of terms: $[R]$ is the set $\{[e] \mid e \in R\}$. By construction, dropping is stronger than lowering, *i.e.* dropping after lowering is the same as dropping, $[[e]_j] = [e]$. As for lowering, dropping commutes with substitution: $[e[x/f]] = [e][x/[f]]$. We overload the notation \mathcal{S} for the set of sound indexed terms. Although it is defined as for lambda terms as $(\rightsquigarrow^* \mathcal{U})$, the meaning is different since the reduction is now bridled by indices.

The calculus on indexed terms is just an instrumentation of the λ -calculus that behaves the same up to the consumption of all the fuel. Formally, we show that they can simulate one another, up to some condition on the indices.

Indexed terms can be simulated by lambda terms. That is, if an indexed term can reduce, then the same reduction step can be performed after dropping indices.

Lemma 74 (Forward simulation). *If $e \rightsquigarrow f$, then $[e] \rightsquigarrow [f]$.*

$$\begin{array}{ll}
P x^k = P k & P \langle e, f \rangle^k = P k \wedge P e \wedge P f \\
P (\lambda^k x e) = P k \wedge P e & P (\text{fst}^k e) = P k \wedge P e \\
P (e f)^k = P k \wedge P e \wedge P f & P (\text{snd}^k e) = P k \wedge P e
\end{array}$$

Figure 5.15: Lifting integer predicates to indexed terms

$$\begin{array}{l}
x^{k_1} \star x^{k_2} = k_1 \star k_2 \\
(\lambda^{k_1} x e_1) \star (\lambda^{k_2} x e_2) = k_1 \star k_2 \wedge e_1 \star e_2 \\
(e_1 f_1)^{k_1} \star (e_2 f_2)^{k_2} = k_1 \star k_2 \wedge e_1 \star e_2 \wedge f_1 \star f_2 \\
\langle e_1, f_1 \rangle^{k_1} \star \langle e_2, f_2 \rangle^{k_2} = k_1 \star k_2 \wedge e_1 \star e_2 \wedge f_1 \star f_2 \\
\text{fst}^{k_1} e_1 \star \text{fst}^{k_2} e_2 = k_1 \star k_2 \wedge e_1 \star e_2 \\
\text{snd}^{k_1} e_1 \star \text{snd}^{k_2} e_2 = k_1 \star k_2 \wedge e_1 \star e_2
\end{array}$$

Figure 5.16: Lifting of a binary predicate \star on indices to terms

Proof. The proof is in the Coq lemma `Llanguage.red_drop`. \square

In order to make the other direction concise, we lift predicates on integers to predicates on indexed terms by requiring the predicate to hold for all indices occurring in the term. For instance, $e > 0$ means that the index of all the nodes of e are greater than zero while $e \leq k$ means that the indices in e are smaller or equal to k . This is formally defined on Figure 5.15 (Coq definition `Flanguage.unary_fuel`).

Indexed terms can simulate lambda terms, provided they have enough fuel. This means that if an indexed term has strictly positive indices and can be reduced after dropping its indices, then the same reduction step can be performed on the indexed term.

Lemma 75 (Backward simulation). *If $e > 0$ and $[e] \rightsquigarrow a'$, then there exists e' such that $e \rightsquigarrow e'$ and $[e'] = a'$.*

Proof. The proof is in the Coq lemma `Llanguage.drop_red_exists`. \square

Using this last lemma, we can show that when a sound indexed term e has indices greater or equal to k , then its erasure can safely do at least k steps. This lemma is crucial to transfer the soundness result in the indexed calculus to the λ -calculus. A lambda term is sound if it is sound for any number of steps.

Lemma 76. *If $e \geq k$ and $e \in \mathcal{S}$ hold, then $[e]$ is sound for all paths of size smaller than k steps.*

Proof. The proof is in the Coq lemma `Lsoundness.term_ge_OK`. \square

5.2.3 Semantic types

In order to define semantic types concisely, it is convenient to have a few helper operations on sets of indexed terms. We first lift binary properties on indices to indexed terms. This is done by asking the two terms to share the same skeleton (they drop on the same lambda term) and the indices of corresponding nodes to be related by the property on indices. A formal definition is given on Figure 5.16 (Coq definition `Flanguage.binary_fuel`).

The *interior* of a set R is the set $R\downarrow$ containing all terms smaller than a term in R , *i.e.* $\{f \mid \exists e \in R, f \leq e\}$ (Coq definition `Fsemantics.Dec`). The *contraction* of a set R is the set $(R\rightsquigarrow)$ of all terms obtained by one-step reduction of a term in R , *i.e.* $\{f \mid \exists e \in R, e \rightsquigarrow f\}$ (Coq definition `Fsemantics.Red`).

A *pretype* is a set of sound terms that contains both its interior and its contraction (Coq definition `Fsemantics.C`). We write \mathcal{C} the set of pretypes.

Definition 77 (Pretypes). $\mathcal{C} \stackrel{\text{def}}{=} \{R \subseteq \mathcal{S} \mid R\downarrow \cup (R\rightsquigarrow) \subseteq R\}$

Notice that the empty set and the set of sound terms \mathcal{S} are pretypes. We restrict pretypes to sound terms because we wish types to be certain forms of pretypes and types to only contain sound terms. The closure of pretypes by interior is just technical: if a property holds for some index, we want it to also hold for smaller indices. The main property of pretypes is to be closed by reduction. Types are pretypes that are also closed by a form of expansion. As a first approximation, sound terms that reduce to a term in a type R should also be in R . However, a type R should still not contain unsound terms even if these reduce to some term in R . Moreover, the meaning of a set of terms R is in essence determined by its set of head normal forms, which we call the kernel of R . We use concatenation for intersection of sets of terms. Hence, the kernel of R is ΔR . A type R needs not to contain every head normal form that reduces to some term in R . Consider for example the term e_0 equal to $\lambda x x$ and one of its expansion is the term e_1 equal to $\lambda x (\lambda y x)(x x)$. The sets $\{e_0\}$ and $\{e_0, e_1\}$ have quite different meanings. The second set has additional requirements for the argument: it should be sound to apply it to itself. Notice that by definition, the kernel is an idempotent operation: $\Delta(\Delta R) = \Delta R$.

The expansion-closure of a set of terms R , written $\diamond R$, is the set $(\rightsquigarrow^*(\nabla \mathcal{U} \uplus \Delta R))$, which contains terms of which every reduction path leads to either a valid neutral term or a head normal form of R . By definition, the expansion closure is monotonic: if $R \subseteq S$, then $\diamond R \subseteq \diamond S$; it is also idempotent: $\diamond(\diamond R) = \diamond R$.

Finally, semantic types are pretypes that are stable by expansion closure (Coq definition `Fsemantics.CE`):

Definition 78 (Semantic types). $\mathbb{T} \stackrel{\text{def}}{=} \{R \in \mathcal{C} \mid \diamond R \subseteq R\}$.

The kernel of a type is a pretype—but not a type. Conversely, the expansion-closure of a pretype is a type. Actually expansion-closure and kernel are almost inverse of one another: if R is a type, then $\diamond(\Delta R) = R$.

The smallest type, called the bottom type and written \perp , is equal to $\diamond\{\}$, that is $(\rightsquigarrow^*(\nabla \mathcal{U}))$. The largest type, \top , called the top type is the set \mathcal{S} of sound terms.

5.2.4 Simple types

We can now define the semantics of arrows and products as semantic type operators. We overload the arrow and product notations, but there is no ambiguity with the syntactic type operators since the semantic operators take semantic types R or S , while syntactic type operators take syntactic types τ or σ as arguments.

Definition 79 (Arrow and product operators).

$$\begin{aligned} R \rightarrow S &\stackrel{\text{def}}{=} \diamond \{ \lambda^k x e \in \mathcal{S} \mid k > 0 \Rightarrow \forall f, [f]_{k-1} \in R \Rightarrow [e[x/f]]_{k-1} \in S \} \\ R \times S &\stackrel{\text{def}}{=} \diamond \{ \langle e, f \rangle^k \in \mathcal{S} \mid k > 0 \Rightarrow [e]_{k-1} \in R \wedge [f]_{k-1} \in S \} \end{aligned}$$

The arrow semantic operator is the expansion closure of sound term abstractions $\lambda^k x e$ satisfying the following property when the index k is non-zero. For all arguments f such that its lowering by $k - 1$ is in the domain \mathbf{R} , the lowering by $k - 1$ of the substitution of x by f in e has to be in the range \mathbf{S} . The product semantic operator is similar. It is the expansion closure of sound pairs satisfying the following property for non-zero indices k . Each component of the pair has to be in its associated type after lowering by $k - 1$.

In order to prove that arrow and product semantic type operators preserve types, we need to define and prove the following easy properties on indices:

- $\llbracket [e]_j \rrbracket_k = [e]_{kj}$ (Coq lemma `Flanguage.lower_lower`)
- If $k' \leq k$ and $e' \leq e$, then $[e']_{k'} \leq [e]_k$. (Coq lemma `Flanguage.le_term_lower`)
- If $e' \leq e$ and $f' \leq f$, then $e'[x/f'] \leq e[x/f]$. (Coq lemma `Flanguage.le_term_subst`)

And this less easy one:

Lemma 80. *If $e \rightsquigarrow f$ holds, then for all k , $[e]_{k+1} \rightsquigarrow f'$ and $[f]_k \leq f'$ hold for some f' .*

Proof. Coq lemma `Flanguage.red_lower` □

The arrow and product operators preserve types.

Lemma 81. *If \mathbf{R} and \mathbf{S} are types, then so are $\mathbf{R} \rightarrow \mathbf{S}$ and $\mathbf{R} \times \mathbf{S}$.*

Proof. Coq lemma `CE_EArr` and `CE_EProd` in `Fsemantics.v`.

We only detail the proof for the arrow operator, which uses indexed terms in a crucial way. The proof for the product operator is similar, but easier. Since the arrow operator is defined by expansion-closure, it is a type if its kernel is a pretype. Its kernel contains only sound terms by definition. So it remains to show that the definition contains its interior and contraction.

Let $\lambda^k x e' \leq \lambda^k x e$ (1), $\lambda^k x e \in \mathcal{S}$ (2), and $k > 0 \Rightarrow \forall f, [f]_{k-1} \in \mathbf{R} \Rightarrow [e[x/f]]_{k-1} \in \mathbf{S}$ (3), and show that $\lambda^k x e' \in \mathcal{S}$ (4) and $j > 0 \Rightarrow \forall f, [f]_{j-1} \in \mathbf{R} \Rightarrow [e'[x/f]]_{j-1} \in \mathbf{S}$ (5). The first assertion (4) comes easily with (1) and (2) since \mathcal{S} contains its interior. To show (5), let $j > 0$ and $[f]_{j-1} \in \mathbf{R}$ (6) and show $[e'[x/f]]_{j-1} \in \mathbf{S}$ (7). By (1) we have $j \leq k$, so $k > 0$. We also have $\llbracket [f]_{j-1} \rrbracket_{k-1} = [f]_{j-1}$ which is in \mathbf{R} by (6). So from (3) we have $[e[x/[f]_{j-1}]]_{k-1} \in \mathbf{S}$. Since \mathbf{S} is a type, it contains its interior so $[e'[x/[f]_{j-1}]]_{j-1} \in \mathbf{S}$. Since the substitution and the lowering function commute, we conclude (7).

Let $\lambda^k x e \rightsquigarrow e_1$ (8), $\lambda^k x e \in \mathcal{S}$ (9), and $k > 0 \Rightarrow \forall f, [f]_{k-1} \in \mathbf{R} \Rightarrow [e[x/f]]_{k-1} \in \mathbf{S}$ (10). By inversion of the reduction relation we have $k = k' + 1$ and $e_1 = \lambda^{k'} x e'$ for some k' and e' such that $e \rightsquigarrow e'$ (11). We now have to show that $\lambda^{k'} x e' \in \mathcal{S}$ (12) and $k' > 0 \Rightarrow \forall f, [f]_{k'-1} \in \mathbf{R} \Rightarrow [e'[x/f]]_{k'-1} \in \mathbf{S}$ (13). We show (12) with (8) and (9) since \mathcal{S} contains its contraction. To show (13), let $k' > 0$ and $[f]_{k'-1} \in \mathbf{R}$ (14) and show $[e'[x/f]]_{k'-1} \in \mathbf{S}$ (15). We have $\llbracket [f]_{k'-1} \rrbracket_{k-1} = [f]_{k'-1}$ which is in \mathbf{R} by (14). So from (10) we have $[e[x/[f]_{k'-1}]]_{k-1} \in \mathbf{S}$. Since \mathbf{S} is a type, it contains its contraction and interior so $[e'[x/[f]_{k'-1}]]_{k'-1} \in \mathbf{S}$ by Lemma 80. Since the substitution and the lowering function commute, we conclude (15). □

5.2.5 Intersection types

The intersection $\bigcap_{i \in I} R_i$ of a nonempty family of types $(R_i)^{i \in I}$ is a type. As a particular case, the bottom type \perp is also the intersection of all types. We extend this definition to the empty family with the top type, which is the set of sound terms. Because all types are sets of sound terms, we can define the intersection operator of a family of types as the intersection with the set of sound terms (Coq definition `Fsemantics.EFor`).

Definition 82 (Intersection operator). *We define $\forall I F$ as the set of terms $\mathcal{S} \cap \bigcap_{i \in I} F i$.*

The intersection operator preserves semantic types:

Lemma 83. *If for all $i \in I$, $F i$ is a semantic type, then $\forall I F$ is a semantic type.*

Proof. Coq lemma `Fsemantics.CE_EFor` □

5.2.6 Recursive types

This section follows the usual description of recursive types using the notion of approximations as done in [4]. Adding recursive types is the main reason for using a step-indexed technique, assuming we have a semantics approach (which is itself due to our notion of coherent abstraction). While recursive types motivate the use of step-indexed semantics, they do not raise any additional difficulty once the semantics has been correctly set up.

The k -approximation of a set R , written $\langle R \rangle_k$ is the subset $\{e \in R \mid e < k\}$ of elements of R that are smaller than k (Coq definition `Fsemantics.approx`). The following properties of approximations immediately follow from the definition: $\langle R \rangle_0$ is the empty set; a sequence of approximations is the approximation by the minimum of the sequence: $\langle \langle R \rangle_j \rangle_k = \langle R \rangle_{jk}$; two sets of terms that are equal at all approximations are equal: if $\langle R \rangle_k = \langle S \rangle_k$ holds for all k , then $R = S$.

Definition 84 (well-foundedness). *A function F on sets of terms is well-founded (resp. non-expansive) if for any set of terms R , the approximations of $F R$ and $F \langle R \rangle_k$ are equal at rank $k + 1$ (resp. k), i.e. $\langle F R \rangle_{k+1} = \langle F \langle R \rangle_k \rangle_{k+1}$ (resp. $\langle F R \rangle_k = \langle F \langle R \rangle_k \rangle_k$)*

Intuitively, well-foundedness (resp. non-expansiveness) ensures that given a term sound for k steps, F returns a term sound for $k+1$ (resp. k) steps. The Coq definition of well-foundedness (resp. non-expansiveness) is given in `WF` (resp. `NE`) in `Fsemantics.v`.

The iteration of a well-founded function F does not look at its argument for terms of small indices: $\langle F^k R \rangle_k$ is independent of R ; in particular, it is equal to $\langle F^k \perp \rangle_k$. This lemma helps us to prove the more useful lemmas that allow us to prove that a term is in the k -th iteration of a well-founded functor, given it is already in the j -th iteration and smaller than both j and k : $\langle F^j R \rangle_{kj}$ and $\langle F^k R \rangle_{kj}$ are equal.

Definition 85 (Recursive operator). *Given a well-founded function F on sets of terms, we define μF as the set of terms $\bigcup_{k \geq 0} \langle F^k \perp \rangle_k$.*

The recursive operator preserves semantic types:

Lemma 86. *If F is well-founded and maps semantic types to semantic types, then μF is a semantic type.*

Proof. Coq lemma `Fsemantics.CE_EMu` □

Moreover, recursive types can be unfolded or folded as expected: if F is well-founded, then $\mu F = F(\mu F)$. This is proved by showing that $\langle \mu F \rangle_k$ is equal to both $\langle F^k \perp \rangle_k$ and $\langle F(\mu F) \rangle_k$ for every k . The proofs are done in Coq lemma `Mu_fold` and `Mu_unfold` in file `Fsemantics.v`.

The following Lemma, although in a different setting, is stated exactly as with traditional step-indexed semantics [4]:

Lemma 87. *We have the following properties:*

- *Every well-founded function is non-expansive.*
- $X \mapsto X$ *is non-expansive.*
- $X \mapsto R$ *where X is unused in R (R is constant) is well-founded.*
- *The composition of non-expansive functors is non-expansive.*
- *The composition of a non-expansive functor with a well-founded functor (in either order) is well-founded.*
- *If F and G are non-expansive, then $X \mapsto (FX \rightarrow GX)$ and $X \mapsto (FX \times GX)$ are well-founded.*
- *If $(F_i)^{i \in I}$ is a family of non-expansive (resp. well-founded) functors, then the functor $X \mapsto \bigcap_{i \in I} (F_i X)$ is non-expansive (resp. well-founded).*
- *If $X \mapsto FXY$ is non-expansive (resp. well-founded) for every Y and FX is well founded for every X , then $X \mapsto \mu(FX)$ is non-expansive (resp. well-founded).*

Proof. Lemma `NE_id`, `WF_CST`, `WF_PArr`, `WF_PProd`, `WFj_For`, and `WFj_Mu` in the `Fsemantics.v` Coq file. □

Just for illustration $X \mapsto X \rightarrow \mathcal{S}$ is well-founded since $X \mapsto X$ is non-expansive and $X \mapsto \mathcal{S}$ is constant, thus well-founded, and therefore non-expansive.

5.2.7 Semantic judgment

A binding is a pair $(x : R)$ of a variable and a semantic type. A context G is a set of bindings, defining a finite mapping from term variables to semantic types. We say that a substitution γ is *compatible* with a context G , written $\gamma : G$, if $\text{dom}(\gamma)$ and $\text{dom}(G)$ coincide and for all $(x : R)$ in G , the term γx is in R .

We define the semantic judgment $G \models S$ as the set of terms e such that γe is in S for any substitution γ “compatible” with G (Coq definition `Fsemantics.EJudg`).

Definition 88 (Semantic judgment).

$$\begin{aligned} \gamma : G &\stackrel{\text{def}}{=} \forall (x : R) \in G, \gamma x \in R \\ G \models S &\stackrel{\text{def}}{=} \{e \mid \forall \gamma : G, \gamma e \in S\} \end{aligned}$$

We may now present the semantic typing rules for the STLC.

Lemma 89 (Variable). *If R is a type and $(x : R)$ is in G , then x^k is in $G \models R$.*

Proof. Coq lemma `Fsemantics.EVar_sem`.

Let γ be compatible with \mathbf{G} (1). We show that γx^k is in \mathbf{R} . Since $(x : \mathbf{R})$ is in \mathbf{G} , we have γx in \mathbf{R} by (1), Being a type, \mathbf{R} is closed by lowering. Hence, $\lfloor \gamma x \rfloor_k$ is also in \mathbf{R} . By definition of substitution, this is equal to γx^k , which is thus also in \mathbf{R} . \square

Lemma 90 (Abstraction). *If \mathbf{R} and \mathbf{S} are types and e is in $\mathbf{G}, (x : \mathbf{R}) \models \mathbf{S}$, then $\lambda^k x e$ is in $\mathbf{G} \models \mathbf{R} \rightarrow \mathbf{S}$.*

Proof. Coq lemma `Fsemantics.ELam_sem`.

Let γ be compatible with \mathbf{G} (1). We show that $\gamma(\lambda^k x e)$ is in $\mathbf{R} \rightarrow \mathbf{S}$ (2). Assume $\gamma(\lambda^k x e) \rightsquigarrow^* e_1$. Then e_1 is necessarily of the form $\lambda^j x e'$ where $\gamma e \rightsquigarrow^* e'$.

We first show that $\lambda^j x e' \in \mathcal{S}$ (3). Since γ is compatible with \mathbf{G} , $\gamma, x \mapsto x$ is compatible with $\mathbf{G}, (x : \mathbf{R})$ as variables are in all types. Since e is in $(\mathbf{G}, (x : \mathbf{R}) \models \mathbf{S})$, we have $(\gamma, x \mapsto x) e$, i.e. γe in \mathbf{S} . Since \mathbf{S} is closed by reduction, we have e' in \mathbf{S} and a fortiori in \mathcal{S} . This implies (3).

Assume $j > 0$ and $\lfloor f \rfloor_{j-1} \in \mathbf{R}$. Let γ' be $\gamma, x \mapsto \lfloor f \rfloor_{j-1}$. By construction $\gamma' : \mathbf{G}, (x : \mathbf{R})$. Since e is in $(\mathbf{G}, (x : \mathbf{R}) \models \mathbf{S})$, we have $\gamma' e$ in \mathbf{S} and, since \mathbf{S} is closed by reduction, $e'[x/\lfloor f \rfloor_{j-1}]$ is also in \mathbf{S} . By decreasing index we have $\lfloor e'[x/\lfloor f \rfloor_{j-1}] \rfloor_{j-1} \in \mathbf{S}$, from which by Lemma 72 becomes $\lfloor e'[x/f] \rfloor_{j-1} \in \mathbf{S}$. This ends the proof of (2). \square

Lemma 91 (Application). *If \mathbf{R} and \mathbf{S} are types, e is in $\mathbf{G} \models \mathbf{R} \rightarrow \mathbf{S}$, and f is in $\mathbf{G} \models \mathbf{R}$, then $(e f)^k$ is in $\mathbf{G} \models \mathbf{S}$ for any k .*

Proof. Coq lemma `Fsemantics.EApp_sem`.

Let γ be compatible with \mathbf{G} . We show that $\gamma(e f)^k \in \mathbf{S}$. By hypotheses we have $\gamma e \in \mathbf{R} \rightarrow \mathbf{S}$ and $\gamma f \in \mathbf{R}$. We prove the more general result that for all k, e , and f , if $e \in \mathbf{R} \rightarrow \mathbf{S}$ and $f \in \mathbf{R}$ hold, then $(e f)^k \in \mathbf{S}$ also holds. This is proved by induction over the strong normalization of e and f using the closure expansion of \mathbf{S} .

The term $(e f)^k$ is neutral. It is also valid since e and f are sound and, by construction of $\mathbf{R} \rightarrow \mathbf{S}$, e is an abstraction when in normal form. If $(e f)^k$ reduces by a context rule, we use our induction hypothesis. Otherwise, e must be of the form $\lambda^{j+1} x e'$ for some j and e' and k be of the form $k' + 1$ and the reduction is $(e f)^k \rightsquigarrow \lfloor e'[x/f] \rfloor_{jk'}$. It remains to show $\lfloor e'[x/f] \rfloor_{jk'} \in \mathbf{S}$ (1). We have $\lfloor f \rfloor_j \in \mathbf{R}$ by stability under decreasing index. So, we have $\lfloor e'[x/f] \rfloor_j \in \mathbf{S}$ by definition of the arrow operator. Then (1) follows by stability under decreasing index. \square

Lemma 92 (Pair). *If \mathbf{R}_i is a type and e_i is in $\mathbf{G} \models \mathbf{R}_i$, then $\langle e_1, e_2 \rangle^k$ is in $\mathbf{G} \models \mathbf{R}_1 \times \mathbf{R}_2$.*

Proof. Coq lemma `Fsemantics.EPair_sem`. \square

Lemma 93 (Projections). *If \mathbf{R} and \mathbf{S} are types and e in $\mathbf{G} \models \mathbf{R} \times \mathbf{S}$, then $\text{fst}^k e$ is in $\mathbf{G} \models \mathbf{R}$ and $\text{snd}^k e$ is in $\mathbf{G} \models \mathbf{S}$.*

Proof. Coq lemma `EFst_sem` and `ESnd_sem` in file `Fsemantics.v`. \square

Note that when \mathbf{R} is a type for all $(x : \mathbf{R}) \in \mathbf{G}$ and \mathbf{S} is a type, then $\mathbf{G} \models \mathbf{S}$ is a pretype.

5.3 Soundness

The soundness proof is not direct. We translate F_{cc} type system from the λ -calculus to a temporary type system on the indexed calculus. We prove soundness for the indexed calculus type system and migrate the result to the λ -calculus type system. The relation between both type systems is that if a lambda term is well-typed then all indexed terms that drop on this lambda term are also well-typed. And reciprocally, if an indexed term is well-typed, then its dropped lambda term is well-typed too. Both directions preserve the typing (the pair of the environment and type). Notice that only the term judgment needs to be changed since it is the only one talking about terms.

Syntactically, the indexed term judgment $e : \Sigma; \Gamma \vdash \tau$ contains the exact same rules as those of the lambda term judgment. However index annotations now appear on the term node we are typing. This annotation has no constraint, which gives us that if a term is typed with annotations it can be typed without and reciprocally if a term is typed without annotations it can be typed with any annotations.

Lemma 94. *The following assertions hold:*

- *If $e : \Sigma; \Gamma \vdash \tau$ holds, then $[e] : \Sigma; \Gamma \vdash \tau$ holds.*
- *If $a : \Sigma; \Gamma \vdash \tau$ holds, then $e : \Sigma; \Gamma \vdash \tau$ holds for all e such that $[e] = a$.*

Proof. Coq lemma `jterm_aux` and `jterm_aux_rev` in file `Lsoundness.v`. □

We show the extraction lemma which is used to prove soundness. It permits to add premises to some typing rules. These premises have to be present for the induction to work, but are simple consequences of all rules taken at once. The premises in question are specifically marked with the condition `mS v` in the definition of `typesystem.jobj` and `Ftypesystem.jterm`. See Section 5.6 for more details about the paper version and soundness version of the type system.

Lemma 95 (Extraction). *The following assertions hold.*

- *If $\Sigma \vdash \kappa$ and $\emptyset \Vdash \Sigma$ hold, then $\Sigma \Vdash \kappa$ holds.*
- *If $\Sigma \vdash \tau : \kappa$ and $\emptyset \Vdash \Sigma$ hold, then $\Sigma \Vdash \kappa$ holds.*
- *If $\Sigma; \Theta_0; \Theta_1 \vdash P$, $\emptyset \Vdash \Sigma$, $\Sigma \Vdash \Theta_0$, and $\Sigma \Vdash \Theta_1$ hold, then $\Sigma \Vdash P$ holds.*
- *If $\Sigma; \Theta_0; \Theta_1 \vdash^* (\Sigma' \vdash \tau') \triangleright \tau$, $\emptyset \Vdash \Sigma$, $\Sigma \Vdash \Theta_0$, and $\Sigma \Vdash \Theta_1$ hold, then $\Sigma \vdash \Sigma'$ holds and if $\Sigma, \Sigma' \vdash \tau' : \star$ also holds, then $\Sigma \vdash \tau : \star$ holds too.*
- *If $\Sigma \vdash \Sigma'$ holds, then $\Sigma \Vdash \Sigma'$ holds.*
- *If $e : \Sigma; \Gamma \vdash \tau$, $\emptyset \Vdash \Sigma$, and $\Sigma \vdash \Gamma$ hold, then $\Sigma \vdash \tau : \star$ holds.*

Proof. Coq lemma `jobj_extra` and `jterm_extra` in file `typesystemextra.v`. The proof is done by induction and uses the usual weakening and substitution lemma. □

To state and prove the soundness of the indexed type system we interpret (syntactic) kinds, types, propositions, and typing environments as sets of semantic objects, semantic objects, semantic propositions, and mappings from type variables to semantic objects respectively. For each judgment, we also give its semantic interpretation as a mathematical assertion. The Coq

$$\begin{aligned}
|\star|_\eta &= \mathbb{T} \\
|1|_\eta &= \{\langle \rangle\} \\
|\kappa_1 \times \kappa_2|_\eta &= \left\{ \langle x_1, x_2 \rangle \mid x_1 \in |\kappa_1|_\eta \wedge x_2 \in |\kappa_2|_\eta \right\} \\
|\{\alpha : \kappa \mid \mathbf{P}\}|_\eta &= \left\{ x \in |\kappa|_\eta \mid \forall k \ |\mathbf{P}|_{\eta, \alpha \mapsto x}^k \right\}
\end{aligned}$$

Figure 5.17: Kind interpretation

$$\begin{aligned}
|\alpha|_\eta &= \eta(\alpha) \\
|\tau \rightarrow \sigma|_\eta &= |\tau|_\eta \rightarrow |\sigma|_\eta & |\langle \rangle|_\eta &= \langle \rangle \\
|\tau \times \sigma|_\eta &= |\tau|_\eta \times |\sigma|_\eta & |\langle \tau, \sigma \rangle|_\eta &= \langle |\tau|_\eta, |\sigma|_\eta \rangle \\
|\forall(\alpha : \kappa) \tau|_\eta &= \forall |\kappa|_\eta (x \mapsto |\tau|_{\eta, \alpha \mapsto x}) & |\mathbf{fst} \tau|_\eta &= \mathbf{fst} |\tau|_\eta \\
|\mu\alpha \tau|_\eta &= \mu (x \mapsto |\tau|_{\eta, \alpha \mapsto x}) & |\mathbf{snd} \tau|_\eta &= \mathbf{snd} |\tau|_\eta \\
|\perp|_\eta &= \perp \\
|\top|_\eta &= \top
\end{aligned}$$

Figure 5.18: Type interpretation

interpretation function is `Fsoundness.semobj`. It is defined as a binary relation, but we show in `semobj_eq` that it behaves as a function.

We define the interpretation of kinds on Figure 5.17. Kinds are interpreted as sets of semantic objects under a mapping η from type variables α to semantic objects. The star kind is interpreted as the set of semantic types \mathbb{T} . The unit kind is interpreted as the singleton set containing the unit object $\langle \rangle$. The product kind $\kappa_1 \times \kappa_2$ is interpreted as the set of pairs of which the first component is in the interpretation of κ_1 and the second component is in the interpretation of κ_2 . Finally the constrained kind $\{\alpha : \kappa \mid \mathbf{P}\}$ is interpreted as the subset of the interpretation of κ for which its elements x satisfy the interpretation of the proposition \mathbf{P} under the extended mapping where α maps to x and for all k .

The interpretation of syntactic types, given in Figure 5.18, are semantic objects, and it is parametrized over a mapping from type variables to semantic objects written η . Semantic objects may be semantic types, the unit object $\langle \rangle$, or pairs of semantic objects $\langle x_1, x_2 \rangle$. The interpretation of a type variable is its value in the mapping. If it is not present in the mapping, the unit semantic object is returned. The interpretation of arrow and product types simply use the arrow and product operators defined in 5.2.4. The interpretation of the polymorphic type $\forall(\alpha : \kappa) \tau$ under η is the intersection of all interpretations of τ under η extended with α mapping to x in the interpretation of κ under η . This is defined using the intersection operator of Section 5.2.5. The interpretation of the recursive type $\mu\alpha \tau$ under η is the infinite iteration of the functor mapping \mathbf{X} under the extension of η mapping α to \mathbf{X} —which corresponds to the infinite unfolding of the recursive type (see Section 5.2.6). Finally, top and bottom are mapped to their semantic equivalent.

The unit type is interpreted as the unit semantic object. The syntactic pair of types τ and σ is interpreted as the semantic pair of the interpretations of τ and σ under the same mapping. Similarly for the projections: `fst` τ and `snd` τ are interpreted to the mathematical projections of respectively the first and second component of semantic pairs.

The interpretation of propositions is given on Figure 5.19. A proposition is interpreted to

$$\begin{aligned}
|\top|_\eta^k &= \top \\
|\mathbf{P}_1 \wedge \mathbf{P}_2|_\eta^k &= |\mathbf{P}_1|_\eta^k \wedge |\mathbf{P}_2|_\eta^k \\
|(\Sigma' \vdash \tau') \triangleright \tau|_\eta^k &= \forall e < k \ (\forall \eta' \in |\Sigma'|_\eta \ e \in |\tau'|_{\eta, \eta'} \Rightarrow e \in |\tau|_\eta) \\
|\exists \kappa|_\eta^k &= \exists x \in |\kappa|_\eta \\
|\forall(\alpha : \kappa) \mathbf{P}|_\eta^k &= \forall x \in |\kappa|_\eta \ |\mathbf{P}|_{\eta, \alpha \mapsto x}^k
\end{aligned}$$

Figure 5.19: Proposition interpretation

a mathematical assertion and depends on a mapping η from type variables to semantic objects and on an index k used to control coinduction. The true and conjunction propositions are interpreted as the true and conjunction mathematical assertions. The interpretation of the coercion proposition $(\Sigma' \vdash \tau') \triangleright \tau$ under η and k asks $\forall |\Sigma'|_\eta$ ($\eta' \mapsto |\tau'|_{\eta, \eta'}$) to be included in $|\tau|_\eta$ for indexed terms e with indices smaller than k . Because all indexed terms are controlled by an index k (Coq lemma `Flanguage.term_lt_exists`), if the inclusion holds for all indices, then it holds for the specific index k , and thus it holds for the term e . The interpretation of the coherence proposition $\exists \kappa$ under η (the coinduction index k is not used) is the assertion that there exists a semantic object x in the interpretation of the kind κ under η . Finally, the polymorphic proposition $\forall(\alpha : \kappa) \mathbf{P}$ is interpreted under η and k by a quantified assertion that \mathbf{P} has to hold for k and the mapping η extended with α mapped to x for all objects x in the interpretation of κ under η .

We define the interpretation of type environments as a set of semantic object mappings (from type variables to semantic objects) in Figure 5.20. The interpretation is parametrized by a surrounding mapping η . The empty environment is interpreted by the singleton set containing the empty mapping \emptyset . The interpretation of an environment Σ extended with a type binding $(\alpha : \kappa)$ extends the mapping η' in the interpretation of Σ under η with α bound to a semantic object in the interpretation of κ under the extended mapping η, η' . We write $|\Sigma|$ to stand for $|\Sigma|_\emptyset$. The following lemma and corollary demonstrate how mapping extension relates to type environment concatenation. If η is a mapping of the interpretation of the type environment Σ and η' is a mapping of the interpretation of Σ' under the surrounding mapping η , then η' is actually an extension of η for the interpretation of the concatenation Σ, Σ' .

Lemma 96. *If $\eta_2 \in |\Sigma_2|_{\eta_1}$ and $\eta_3 \in |\Sigma_3|_{\eta_1, \eta_2}$, then $\eta_2, \eta_3 \in |\Sigma_2, \Sigma_3|_{\eta_1}$.*

Corollary 97. *If $\eta \in |\Sigma|$ and $\eta' \in |\Sigma'|_\eta$, then $\eta, \eta' \in |\Sigma, \Sigma'|$.*

Proof. Coq lemma `Fsoundness.Happ_fH` □

The interpretation of proof environments is given in Figure 5.20. The interpretation $|\Theta|_\eta^k$ of the proof environment Θ under the mapping η and with index k is a mathematical proposition. The interpretation of the empty environment is the tautology proposition, which is always true. The extension of an environment Θ with proposition \mathbf{P} under η with k is the conjunction of the interpretations of Θ and \mathbf{P} under the same arguments η and k . Said otherwise $|\Theta|_\eta^k$ holds if and only if all its propositions hold under η and k .

Finally, the interpretation of term environments is given in Figure 5.20. The interpretation $|\Gamma|_\eta$ of the term environment Γ under the mapping η is the semantic context \mathbf{G} where all syntactic types have been interpreted to a semantic type \mathbf{R} under η .

$$\begin{aligned}
|\emptyset|_\eta &= \{\emptyset\} \\
|\Sigma, (\alpha : \kappa)|_\eta &= \left\{ \eta', \alpha \mapsto x \mid \eta' \in |\Sigma|_\eta \wedge x \in |\kappa|_{\eta, \eta'} \right\} \\
|\emptyset|_\eta^k &= \top \\
|\Theta, \mathbf{P}|_\eta^k &= |\Theta|_\eta^k \wedge |\mathbf{P}|_\eta^k \\
|\emptyset|_\eta &= \emptyset \\
|\Gamma, (x : \tau)|_\eta &= |\Gamma|_\eta, (x : |\tau|_\eta)
\end{aligned}$$

Figure 5.20: Environments interpretation

The semantic weakening lemma tells that the interpretation of a syntactical object (a kind, a type, a proposition, a type environment, a proof environment, or a term environment) under the mapping η , only looks at the values of η for its free type variables. As a consequence, we may freely change the mapping η as long as we do not touch the semantic objects associated to the free type variables of the syntactical object we consider, it will not change its interpretation.

Lemma 98 (Semantic weakening). *If η' and η agree on the free variables of κ (resp. τ , \mathbf{P} , Σ , and Θ), then $|\kappa|_{\eta'} = |\kappa|_\eta$ (resp. $|\tau|_{\eta'} = |\tau|_\eta$, $|\mathbf{P}|_{\eta'}^k = |\mathbf{P}|_\eta^k$ for all k , $|\Sigma|_{\eta'} = |\Sigma|_\eta$, $|\Theta|_{\eta'}^k = |\Theta|_\eta^k$ for all k , and $|\Gamma|_{\eta'} = |\Gamma|_\eta$) holds.*

Proof. Coq lemma `Fsoundness.semobj_lift` □

The semantic substitution lemma relates type substitution with mapping extensions. The interpretation under η of a semantic object in which we substituted the type variable α with type σ is equal to the interpretation of the same object without substitution under the extension of η where α maps to the interpretation of σ under η .

Lemma 99 (Semantic substitution). *Let $\eta' = \eta, \alpha \mapsto |\sigma|_\eta$. The following assertions hold.*

- $|\kappa[\alpha/\sigma]|_\eta = |\kappa|_{\eta'}$ holds.
- $|\tau[\alpha/\sigma]|_\eta = |\tau|_{\eta'}$ holds.
- $|\mathbf{P}[\alpha/\sigma]|_\eta^k = |\mathbf{P}|_{\eta'}^k$ holds for all k .
- $|\Sigma[\alpha/\sigma]|_\eta = |\Sigma|_{\eta'}$ holds.
- $|\Theta[\alpha/\sigma]|_\eta^k = |\Theta|_{\eta'}^k$ holds for all k .
- $|\Gamma[\alpha/\sigma]|_\eta = |\Gamma|_{\eta'}$ holds.

Proof. Coq lemma `Fsoundness.semobj_subst` □

We can now define and prove the soundness of each syntactic judgment according to its semantic. If τ is non-expansive (resp. well-founded) with respect to α , then its interpretation as a functor under any mapping is non-expansive (resp. well-founded). If a kind κ is coherent under environment Σ , then its interpretation under any mapping η of the interpretation of Σ is inhabited. If a type τ has kind κ under the environment Σ , then its interpretation is in the interpretation of its kind, for all mappings η of the interpretation of Σ .

$$\begin{array}{ll}
[x]^k = x^k & \langle [a, b] \rangle^k = \langle [a]^k, [b]^k \rangle^k \\
[\lambda x a]^k = \lambda^k x [a]^k & [\text{fst } a]^k = \text{fst}^k [a]^k \\
[ab]^k = ([a]^k [b]^k)^k & [\text{snd } a]^k = \text{snd}^k [a]^k
\end{array}$$

Figure 5.21: Fill function

If the proposition P holds under Σ , Θ_0 , and Θ_1 , then for all mappings η of the interpretation of Σ and for all indices k , if the interpretation of Θ_0 holds under η for all indices j smaller than or equal to k and the interpretation of Θ_1 holds under η for all indices j smaller than k , then the interpretation of P holds under η with k . We see in the interpretation of the proposition judgment that Θ_0 can be used at level k while Θ_1 is only accessible at level $k - 1$. So in order for P to use an hypothesis in Θ_1 , it has to do something constructive in order to decrease the level and use the hypothesis.

If a type environment Σ' is coherent under Σ , then its interpretation is inhabited under any mapping η of the interpretation of Σ . If the term environment Γ is valid under the type environment Σ , then for all mappings η of the interpretation of Σ the interpretation under η of the types of Γ are semantic types. If the indexed term e has type τ under Σ and Γ , then for all type mappings η of the interpretation of Σ the term e is in the semantic judgment $|\Gamma|_\eta \models |\tau|_\eta$.

Lemma 100 (Judgment soundness). *The following assertions hold.*

- If $\alpha \mapsto \tau : \text{NE}$ holds, then $\forall \eta \ X \mapsto |\tau|_{\eta, \alpha \mapsto X}$ is non-expansive.
- If $\alpha \mapsto \tau : \text{WF}$ holds, then $\forall \eta \ X \mapsto |\tau|_{\eta, \alpha \mapsto X}$ is well-founded.
- If $\Sigma \vdash \kappa$ holds, then $\forall \eta \in |\Sigma| \ |\kappa|_\eta \neq \emptyset$.
- If $\Sigma \vdash \tau : \kappa$ holds, then $\forall \eta \in |\Sigma| \ |\tau|_\eta \in |\kappa|_\eta$.
- If $\Sigma; \Theta_0; \Theta_1 \vdash P$ holds, then $\forall \eta \in |\Sigma| \ \forall k \ (\forall j \leq k \ |\Theta_0|_\eta^j) \wedge (\forall j < k \ |\Theta_1|_\eta^j) \Rightarrow |P|_\eta^k$.
- If $\Sigma \vdash \Sigma'$ holds, then $\forall \eta \in |\Sigma| \ |\Sigma'|_\eta \neq \emptyset$.
- If $\Sigma \vdash \Gamma$ holds, then $\forall \eta \in |\Sigma| \ \forall (x : R) \in |\Gamma|_\eta \ R \in \mathbb{T}$.
- If $e : \Sigma; \Gamma \vdash \tau$ holds, then $\forall \eta \in |\Sigma| \ e \in |\Gamma|_\eta \models |\tau|_\eta$.

Proof. Lemma `jrec_sound`, `jobj_sound`, and `jterm_sound` in the `Fsoundness.v` Coq file. \square

To prove the soundness of System F_{cc} , we need to define the filling $[a]^k$ of a lambda term a at rank k as the indexed term obtained by annotating each node of a with index k (Figure 5.21 and Coq definition `Llanguage.kfill`). By construction, we have $[[a]^k] = a$. The drop and fill functions are used to go back and forth between both type systems.

We can now define and prove the soundness lemma. If a is well-typed under coherent and valid environments then a is sound.

Theorem 101 (Soundness of System F_{cc}). *If $\emptyset \vdash \Sigma$, $\Sigma \vdash \Gamma$, and $a : \Sigma; \Gamma \vdash \tau$ hold, then $a \in \mathcal{S}$.*

Proof. Coq lemma `Lsoundness.soundness`.

The lambda term a is sound if it is sound for any number of steps k . Let's show that it is sound for at least k steps. Let e be $\lceil a \rceil^k$. We have $e \geq k$, and by Lemma 76 it suffices to show that e is sound. We have $e : \Sigma; \Gamma \vdash \tau$ by Lemma 94. By Lemma 95, we have $\emptyset \Vdash \Sigma$ in a first step and $\Sigma \vdash \tau : \star$ in a second step. By Lemma 100, $|\Sigma|$ is inhabited and there is a mapping η such that $\eta \in |\Sigma|$. We also have that all sets in $|\Gamma|_\eta$ are semantic types. And we also have that $|\tau|_\eta$ is a semantic type. And finally we have e is in the semantic judgment $|\Gamma|_\eta \models |\tau|_\eta$, which is a pretype. We deduce $e \in \mathcal{S}$. \square

Termination in the absence of recursive types Although evaluation may not terminate because of the presence of recursive types, it remains interesting to show that recursive types are the only source of non-termination. We already know this for System F. We show that coercions do not themselves introduce non-termination, as long as all types remain non-recursive. The proof is based on reducibility candidates as for System F and does not raise any difficulties. It is almost a copy-paste of the soundness proof omitting step-indices. We thus omit the details.

Theorem 102 (Termination). *If $\emptyset \vdash \Sigma$, $\Sigma \vdash \Gamma$, and $a : \Sigma; \Gamma \vdash \tau$ hold in the sublanguage without recursive types, then a strongly normalizes.*

Proof. Coq lemma `Lnormalization.normalization`. \square

5.4 Expressivity

System F_{cc} goal was to extend System F_t^p to gain the expressiveness of Constraint ML. In order to do so we added both recursive types and coherent coercion abstraction. Recursive types are necessary because consistency in Constraint ML is only defined with recursive types. There is apparently no simple criteria for consistency implying strong normalization in addition to soundness. In System F_{cc} however, it is simple to remove recursive types and thus allow us to write programs with mechanisms similar to Constraint ML which are sound and terminating.

Coherent coercion abstraction was actually added using the coherent polymorphism type system feature. Polymorphism permits to abstract over an erasable content as long as it is coherent. We classify erasable contents, *i.e.* types, with kinds. We added a particular kind which permits to restrain an existing kind for types satisfying a proposition, and propositions contain coercions. In a first section, we define some surface notations. In the next two sections we show how System F_t^p and Constraint ML are included in System F_{cc} . We then discuss the differences between equi-recursive types which we provide in System F_{cc} and iso-recursive types which we defined in System F_{rec} in Section 3.3. The final section shows how our notion of coinduction subsumes existing notions about recursive type equality and recursive type congruence.

5.4.1 Surface notations

To ease readability, we define a few surface notations that easily desugar to System F_{cc} .

Type bindings with constrained kinds The first notation is about type bindings $(\alpha : \kappa)$ when the kind κ is a constrained kind $\{\alpha' : \kappa' \mid P'\}$. The binding is thus written $(\alpha : \{\alpha' :$

$\kappa' \mid P'\}$). On the one hand, the type variable α does not bind in the constrained kind and is thus not free in κ' and P' . But it binds for what comes next (which is the role of a binder); if the binder is used for a polymorphic type $\forall(\alpha : \{\alpha' : \kappa' \mid P'\})\sigma$, it binds in σ ; if it is used in a type environment $\Sigma_1, (\alpha : \{\alpha' : \kappa' \mid P'\}), \Sigma_2$, it binds in Σ_2 and what comes after (coinductive hypotheses or what is after the turnstile). On the other hand, the type variable α' binds in P' , but not in what comes after. It also does bind in its kind κ' . Since these scopes are not ambiguous we may reuse the same type variable for α and α' , and thus write $(\alpha : \{\alpha : \kappa' \mid P'[\alpha'/\alpha]\})$ (which is what happens when using de Bruijn indices).

However these two type variables are essentially the same: they represent the same type. We may thus simply write $(\alpha : \kappa \mid P)$ instead of $(\alpha : \{\alpha : \kappa \mid P\})$. In other words the constrained kind, with parentheses instead of curly braces, becomes a binder for the constrained kind itself, reusing its type variable for the binding type variable. This notation applies only when the constrained kind is used where a binder is expect. Here follow two examples: one for the polymorphic type and one for type environment extension.

$$\begin{aligned} \forall(\alpha : \kappa \mid P)\sigma &\stackrel{\text{def}}{=} \forall(\alpha : \{\alpha : \kappa \mid P\})\sigma \\ \Sigma, (\alpha : \kappa \mid P) &\stackrel{\text{def}}{=} \Sigma, (\alpha : \{\alpha : \kappa \mid P\}) \end{aligned}$$

Existential propositions In type systems with constraints, like **Constraint ML**, there is usually an existential proposition. As we saw in rule **PROPEXI**, our coherence proposition $\exists\kappa$ behaves like an existential. It is introduced with a rule forgetting the witness type, and eliminated with rules **TYPESTR** and **PROPSTR** which are similar to existential elimination. Thus it is no surprise that we can easily encode existential propositions with the coherence proposition of a constrained kind.

$$\exists(\alpha : \kappa)P \stackrel{\text{def}}{=} \exists\{\alpha : \kappa \mid P\}$$

The existential proposition $\exists(\alpha : \kappa)P$ is thus true, if there is a witness type τ in the constrained kind $\{\alpha : \kappa \mid P\}$, which by unfolding rule **TYPEPACK** means that τ has kind κ and satisfies the proposition P .

Lists of type variables and lists of kinds The reason we added the unit kind and product kinds was to simplify the Coq development when dealing with lists of types. Our prior version had a polymorphic type of the form $\forall(\bar{\alpha}, \bar{\tau}_1 \triangleright \bar{\tau}_2)\sigma$, which meant that we abstracted over all type variables in $\bar{\alpha}$ satisfying all the coercions in $\bar{\tau}_1 \triangleright \bar{\tau}_2$. We eliminated these two uses of lists using the unit kind and product kind for the type variables, and the true proposition and conjunction proposition for the coercions. However, a notation of the form $\{\bar{\alpha} : \bar{\kappa} \mid P\}$, and thus $\forall(\bar{\alpha} : \bar{\kappa} \mid P)\sigma$ using the type binding with constrained kind notation, may be useful.

The definition of such notation is not local to the binder, but needs access to what is in the scope of the binder. In the case of the constrained kind $\{\bar{\alpha} : \bar{\kappa} \mid P\}$ we need to modify the proposition P . Similarly, for the polymorphic type $\forall(\bar{\alpha} : \bar{\kappa})\sigma$, we need access to the type σ . And finally, for the combination of both $\forall(\bar{\alpha} : \bar{\kappa} \mid P)\sigma$, we need to modify both P and σ .

We will study the situation with the constrained kind polymorphic type, but it applies for other binding structures. We define the notation $\forall(\bar{\alpha} : \bar{\kappa} \mid P)\sigma$ as $\forall(\alpha' : \kappa' \mid P')\sigma'$, where κ' , P' , and σ' are defined as follows. The type variable α' is taken fresh. We define κ' as the list $\bar{\kappa}$ encoded using the product kind for cons and the unit kind for nil. For example, the list κ_1, κ_2 is encoded as $\kappa_1 \times (\kappa_2 \times 1)$. The proposition P' and the type σ' are

substitutions of P and σ respectively. We substitute each use of α_i (the i th element of the list $\bar{\alpha}$) as the first type projection of $(i - 1)$ iterations of the second type projection of the type variable α' . For example, if the list of type variables was α_1, α_2 , the substitution will be $[\alpha_1/\text{fst } \alpha'][\alpha_2/\text{fst } (\text{snd } \alpha')]$. When instantiating such kind of notations, we need to define the encoding of a list of types. The idea is similar to the encoding of the list of kinds. The list of types $\bar{\tau}$ is encoded as a list where the pair type is used for cons and the unit type is used for nil. For example, the list τ_1, τ_2 is encoded as $\langle \tau_1, \langle \tau_2, \langle \rangle \rangle \rangle$.

Here is a list of other places where this notation may be used: polymorphic types $\forall(\bar{\alpha} : \bar{\kappa}) \sigma$, constrained kinds $\{\bar{\alpha} : \bar{\kappa} \mid P\}$, constrained kind polymorphic types $\forall(\bar{\alpha} : \bar{\kappa} \mid P) \sigma$ as we saw, existential propositions $\exists(\bar{\alpha} : \bar{\kappa}) P$ because they use the constrained kind, type extensions $\Sigma, (\bar{\alpha} : \bar{\kappa})$, and constrained kind type extensions $\Sigma, (\bar{\alpha} : \bar{\kappa} \mid P)$.

The star kind A very simple notation, but easing readability, is to omit the star kinds in type bindings. We may for example write: $\forall \alpha \tau$ for polymorphic types of the kind star, $\forall \bar{\alpha} \tau$ for polymorphic types over a list of types of the kind star, $\forall(\bar{\alpha} \mid P) \sigma$ for polymorphic types over a list of types of the kind star satisfying the proposition P , or $\exists \bar{\alpha} P$ for existential propositions over a list of type variables of the kind star satisfying the proposition P .

Erasable isomorphisms It is sometime useful to have a notion of isomorphisms, like the equality coercions in FC [35]. We can encode this notion using two inverse coercions. We define the isomorphism proposition notation $\tau \equiv \sigma$ as $(\emptyset \vdash \tau) \triangleright \sigma \wedge (\emptyset \vdash \sigma) \triangleright \tau$. However, this definition can only be used with types of kind star, because coercions are only defined at kind star. In order to define a notion of isomorphism at any kind, we lift our notion of erasable isomorphism at kind star to other kinds by extensionality. We index the notation $\tau \equiv_{\kappa} \sigma$ with the kind κ of types τ and σ . We could consider heterogeneous isomorphisms, but we restrict to homogeneous isomorphisms for simplicity. The isomorphism at kind κ is defined inductively on the kind as follows (the last line concerns type-level functions defined in Section 6.1.3):

$$\begin{aligned}
\tau \equiv_{\star} \sigma & \stackrel{\text{def}}{=} \tau \equiv \sigma \\
\tau \equiv_1 \sigma & \stackrel{\text{def}}{=} \top \\
\tau \equiv_{\kappa_1 \times \kappa_2} \sigma & \stackrel{\text{def}}{=} (\text{fst } \tau \equiv_{\kappa_1} \text{fst } \sigma) \wedge (\text{snd } \tau \equiv_{\kappa_2} \text{snd } \sigma) \\
\tau \equiv_{\{\alpha : \kappa \mid P\}} \sigma & \stackrel{\text{def}}{=} \tau \equiv_{\kappa} \sigma \\
\tau \equiv_{\kappa_1 \rightarrow \kappa_2} \sigma & \stackrel{\text{def}}{=} \forall(\alpha \beta : \kappa_1 \times \kappa_1 \mid \alpha \equiv_{\kappa_1} \beta) \tau \alpha \equiv_{\kappa_2} \sigma \beta
\end{aligned}$$

5.4.2 System F_l^P

There is no surprise that System F_{cc} extends System F_l^P . The encoding is simple and not very interesting. And the proof is not difficult but it has to be done. We define two translations from System F_l^P to System F_{cc} : one for types and one for environments. We write $\hat{\tau}$ for the translation of the type τ . We give its definition in Figure 5.22. Type variables, arrows, products, polymorphic types, top and bottom are translated to their equivalent in System F_{cc} . Notice that the polymorphic type uses the notation omitting to mention the kind star (see Section 5.4.1). The bounded polymorphic types are translated to polymorphic types over a constrained kind (we can use the notations we defined in Section 5.4.1). The upper bounded polymorphic type $\forall(\alpha \triangleright \tau) \sigma$ is translated to a polymorphic type with body $\hat{\sigma}$ over the type

$$\begin{array}{l}
\widehat{\alpha} = \alpha \\
\widehat{\tau \rightarrow \sigma} = \widehat{\tau} \rightarrow \widehat{\sigma} \\
\widehat{\tau \times \sigma} = \widehat{\tau} \times \widehat{\sigma} \\
\widehat{\forall \alpha \tau} = \forall \alpha \widehat{\tau} \\
\widehat{\top} = \top \\
\widehat{\perp} = \perp
\end{array}
\quad
\begin{array}{l}
\forall(\widehat{\alpha \triangleright \tau}) \sigma = \forall(\alpha \mid (\emptyset \vdash \alpha) \triangleright \widehat{\tau}) \widehat{\sigma} \\
\forall(\widehat{\alpha \triangleleft \tau}) \sigma = \forall(\alpha \mid (\emptyset \vdash \widehat{\tau}) \triangleright \alpha) \widehat{\sigma}
\end{array}$$

Figure 5.22: System F_t^p type translation function

$$\begin{array}{l}
\widehat{\emptyset}^\Sigma = \emptyset \\
\Gamma, \widehat{(x : \tau)}^\Sigma = \widehat{\Gamma}^\Sigma \\
\widehat{\Gamma, \alpha}^\Sigma = \widehat{\Gamma}^\Sigma, \alpha \\
\Gamma, \widehat{(\tau \triangleright \sigma)}^\Sigma = \widehat{\Gamma}^\Sigma, (\alpha \mid (\emptyset \vdash \widehat{\tau}) \triangleright \widehat{\sigma})
\end{array}
\quad
\begin{array}{l}
\widehat{\emptyset}^\Gamma = \emptyset \\
\Gamma, \widehat{(x : \tau)}^\Gamma = \widehat{\Gamma}^\Gamma, (x : \widehat{\tau}) \\
\widehat{\Gamma, \alpha}^\Gamma = \widehat{\Gamma}^\Gamma \\
\Gamma, \widehat{(\tau \triangleright \sigma)}^\Gamma = \widehat{\Gamma}^\Gamma
\end{array}$$

Figure 5.23: System F_t^p environment translation function

variable α ranging over the constrained kind $\{\alpha \mid (\emptyset \vdash \alpha) \triangleright \tau\}$. This constrained kind contains types (of the star kind) satisfying the proposition that α can be coerced to τ , which is exactly what upper bounded polymorphism requires. Lower bounded polymorphism is similar.

We define the translation for environments in Figure 5.23. Environments can be translated as type environments, written $\widehat{\Gamma}^\Sigma$, or as term environments, written $\widehat{\Gamma}^\Gamma$. The exponent symbol is not a metavariable but an annotation to differentiate the translation. The translation as term environments simply filters the term bindings and translates their types. The translation for type environments translates the sole type bindings to type bindings of the star kind, and bounded type bindings to type bindings of a constrained kind with the coercion as the constraining proposition. Again, we use the notations defined in Section 5.4.1.

To prove the inclusion of System F_t^p in System F_{cc} , we need to prove that if a judgment holds in System F_t^p , then its translation holds in System F_{cc} . We write $(J)_s$ for System F_t^p judgments and $(J)_t$ for System F_{cc} judgments. From the environment judgment we actually extract two translated judgments: one for the type environment and one for the term environment which has to hold under the translated type environment. All types in System F_t^p have the star kind as we can see in the second assertion. And all coercions are typed without coinduction hypotheses as we can see in the third assertion.

Lemma 103. *The following assertions hold.*

- If $(\Gamma \text{ env})_s$ holds, then $(\emptyset \vdash \widehat{\Gamma}^\Sigma)_t$ and $(\widehat{\Gamma}^\Sigma \vdash \widehat{\Gamma}^\Gamma)_t$ hold.
- If $(\Gamma \vdash \tau \text{ type})_s$ holds, then $(\widehat{\Gamma}^\Sigma \vdash \widehat{\tau} : \star)_t$ holds.
- If $(\Gamma \vdash (\Sigma \vdash \tau) \triangleright \sigma)_s$ and $(\Gamma, \Sigma \vdash \tau \text{ type})_s$ hold, then $(\widehat{\Gamma}^\Sigma; \emptyset; \emptyset \vdash (\widehat{\Sigma}^\Sigma \vdash \widehat{\tau}) \triangleright \widehat{\sigma})_t$ holds.
- If $(a : \Gamma \vdash \tau)_s$ holds, then $(a : \widehat{\Gamma}^\Sigma; \widehat{\Gamma}^\Gamma \vdash \widehat{\tau})_t$ holds.

Proof. We proceed by mutual induction. For rule ENVEMPTY, the assertion holds because the empty type environment is well formed, and the empty term environment is well-formed under

well-formed type environment. Rules `ENVTERM` and `ENVTYPE` hold by induction hypotheses without difficulties.

Rule `ENVTYPEL` is more involved. This is actually the most interesting part of the proof because this is where we forge (very simple) witnesses. If we write Σ the translated type environment by induction hypothesis, then we have to show that $\Sigma \vdash \{\alpha \mid (\emptyset \vdash \hat{\tau}) \triangleright \alpha\}$ holds. By rule `KIND` we have to show $\Sigma; \emptyset; \emptyset \vdash \exists \alpha (\emptyset \vdash \hat{\tau}) \triangleright \alpha$. We use rule `TYPEPACK` with the top type. The well-formedness premise for the proposition $\Sigma, \alpha \Vdash (\emptyset \vdash \hat{\tau}) \triangleright \alpha$ uses the induction hypothesis to show $\Sigma, \alpha \vdash \hat{\tau} : \star$. The top type has kind star by rule `TYPETOP`. We prove the last premise $\Sigma; \emptyset; \emptyset \vdash (\emptyset \vdash \hat{\tau}[\alpha/\top]) \triangleright \top$ by rule `COERTOP` and the substitution lemma to show $\Sigma \vdash \hat{\tau}[\alpha/\top] : \star$. Rule `ENVTYPEU` is similar with the bottom type as witness.

The second assertion is not difficult and uses only induction hypotheses. For the third assertion, proof environments are always empty. Rules `COERREFL`, `COERTRANS`, `COERWEAK`, `COERETAARR`, `COERETAPROD`, `COERTOP`, and `COERBOT` have analog rules.

Rule `COERVAR` is interesting because it does not use a coercion variable rule, as there is no such rule. It uses rule `PROPRES` with the instance type equal to the type variable associated to the coercion, which does exist because the environment is well-formed. For instance, if the coercion judgment we have to translate is $\Gamma, \alpha, (\alpha \triangleright \tau), \Gamma' \vdash \alpha \triangleright \tau$ then we have $(\alpha : \{\alpha \mid (\emptyset \vdash \alpha) \triangleright \hat{\tau}\})$ in our translated type environment. With rule `TYPEVAR` we derive $\Sigma \vdash \alpha : \{\alpha \mid (\emptyset \vdash \alpha) \triangleright \hat{\tau}\}$. From which using rule `PROPRES` we get $\Sigma; \emptyset; \emptyset \vdash (\emptyset \vdash \alpha) \triangleright \hat{\tau}$.

Rules `COERTLAM` and `COERTAPP` use rules `COERGEN` and `COERINST` with the star kind. Rules `COERTLAML`, `COERTLAMU`, `COERTAPPL`, and `COERTAPPU` are similar with more elaborated kinds. We detail rule `COERTAPPU`. We have to show that $\Sigma \vdash \hat{\sigma} : \{\alpha \mid (\emptyset \vdash \alpha) \triangleright \hat{\tau}\}$ holds, if we write Σ the translated type environment. We use rule `TYPEPACK`. The first premise $\Sigma \vdash \hat{\sigma} : \star$ is proved by the first induction hypothesis. The second premise $\Sigma; \emptyset; \emptyset \vdash (\emptyset \vdash \hat{\sigma}) \triangleright \hat{\tau}[\alpha/\hat{\sigma}]$ is proved by the second induction hypothesis.

The last assertion only contains analogous rules. \square

5.4.3 Constraint ML

In **Constraint ML**, the main feature is constraint generalization. We write $\hat{\tau}$ the type translation of τ . We can translate constraints as propositions. The constraint $\tau \triangleright \sigma$ becomes the coercion $(\emptyset \vdash \hat{\tau}) \triangleright \hat{\sigma}$. Constraint conjunctions $C_1 \wedge C_2$ are translated to proposition conjunctions $\widehat{C}_1 \wedge \widehat{C}_2$. The existential constraints $\exists \bar{\alpha}. C$ become $\exists \bar{\alpha} \widehat{C}$, which are existential propositions over a series of type variables (see Section 5.4.1). We translate types in the same manner. In particular, type schemes $\forall \bar{\alpha}. C \Rightarrow \sigma$ are translated to $\forall (\bar{\alpha} \mid \widehat{C}) \widehat{\sigma}$. Environments are simply translated by translating their type schemes.

We translate the term judgment $C; \Gamma \vdash a : \tau$ into $a : \emptyset, (\bar{\alpha} \mid \widehat{C}); \widehat{\Gamma} \vdash \hat{\tau}$. The type environment contains a unique type binding for all the type variables (of the star kind using notations of Section 5.4.1) in C , Γ , and τ . The type binding has a constrained kind resulting from the constraint of the **Constraint ML** judgment. Most term typing rules are simple translations. We focus on the most interesting ones.

$$\frac{\text{TERMSUB} \quad C; \Gamma \vdash a : \tau \quad C \vdash \tau \triangleright \tau'}{C; \Gamma \vdash a : \tau'}$$

Rule `TERMSUB` (reminded above) uses rule `TERMCOER` as long as $C \vdash \tau \triangleright \tau'$ can be translated

to a derivation of $\varnothing, (\bar{\alpha} \mid \widehat{C}); \varnothing; \varnothing \vdash (\varnothing \vdash \widehat{\tau}) \triangleright \widehat{\tau}$ which depends on the power of the **Constraint ML** constraint judgment.

We remind rule **TERMINTRO**:

$$\frac{\text{TERMINTRO} \quad C \wedge D; \Gamma \vdash a : \tau \quad \bar{\alpha} \notin \text{fv}(C, \Gamma)}{C \wedge \exists \bar{\alpha}. D; \Gamma \vdash a : \forall \bar{\alpha}. D \Rightarrow \tau}$$

We use $\bar{\beta}$ to designate the type variables of the judgment that are not in $\bar{\alpha}$. By induction hypothesis, we have $a : \varnothing, (\bar{\beta}\bar{\alpha} : \{\bar{\beta}\bar{\alpha} : \bar{\kappa} \times \bar{\kappa} \mid \widehat{C}\theta \wedge \widehat{D}\theta\}); \widehat{\Gamma} \vdash \widehat{\tau}\theta$, where θ is the substitution $[\bar{\alpha}/\text{snd } \bar{\beta}\bar{\alpha}][\bar{\beta}/\text{fst } \bar{\beta}\bar{\alpha}]$. We now weaken the type environment to $\varnothing, (\bar{\beta} \mid \widehat{C} \wedge \exists \bar{\alpha} \widehat{D}), (\bar{\alpha} \mid \widehat{D})$. We clearly have the same type variables with the same kinds (the star kind) because we only partitioned the existing binding. It remains to show that this partition is still coherent. The first binding is coherent by hypothesis. And the second binding $(\bar{\alpha} \mid \widehat{D})$ is coherent under $\varnothing, (\bar{\beta} \mid \widehat{C} \wedge \exists \bar{\alpha} \widehat{D})$ by rules **KIND**, **PROPSND**, **PROPRES**, and **TYPEVAR**.

Rule **TERMELEM** relies on rule **TERMCOER** with **COERINST**. We have to show that $\varnothing, (\bar{\beta} \mid \widehat{C}) \vdash \widehat{\tau} : \{\bar{\alpha} \mid \widehat{D}\}$ holds. Rule **TYPEPACK** requires $\varnothing, (\bar{\beta} \mid \widehat{C}); \varnothing; \varnothing \vdash \widehat{D}[\bar{\alpha}/\widehat{\tau}]$. Depending on the power of **Constraint ML** constraint judgment, we may have a translation from $C \vdash D[\bar{\alpha} \leftarrow \tau]$ to our judgment.

To conclude, System F_{cc} features **Constraint ML** mechanisms for the term judgment. Since **Constraint ML** is parametrized over its constraint mechanism, the inclusion is complete if the constraint judgment implies its translation as a System F_{cc} proposition. In particular, constraint implication $C \vdash D$ should imply its translation, which is the proposition $\varnothing, (\bar{\alpha} \mid \widehat{C}); \varnothing; \varnothing \vdash \widehat{D}$ where $\bar{\alpha}$ are the free variables of the constraints C and D . A simple situation when this implication holds is when the constraint mechanism exhibits type witnesses. These type witnesses can then be used to build a derivation of the translated proposition.

5.4.4 Recursive coercions

We added recursive types to System F_{cc} , because the solvability condition of **Constraint ML** uses them ([28] and [31]). There are two sorts of recursive types: equi-recursive types and iso-recursive types. Equi-recursive types leave no traces in terms. They define the unfolding of recursive types as a type equality. The types $\mu\alpha\tau$ and $\tau[\alpha/\mu\alpha\tau]$ are really the same type. By contrast, these are different types in the iso-recursive view and fold and unfold coercions allow to change one into the other. The version of System F_{rec} we presented in Section 3.3 uses iso-recursive types.

Equi-recursive types are also stronger than iso-recursive types for two reasons. First, because the folding and unfolding of equi-recursive types may be used at any depth. Notice however, that in type systems with subtyping or coercions, this is not a difference. Then, because they may equate two recursive types of different periods. For instance, the types $\mu\alpha \text{Int} \rightarrow \alpha$ and $\mu\alpha \text{Int} \rightarrow \text{Int} \rightarrow \alpha$ are equal with equi-recursive types, but there is no possible conversion from one to the other with iso-recursive types.

In System F_{cc} , although recursive types are folded and unfolded in our coercion judgment (as it is the case with iso-recursive types) and not in the type equality (as it is the case with equi-recursive types), they have the expressivity of equi-recursive types. This is possible for two reasons. First, our recursive types do not leave any trace at the term level: they are erasable. Erasability is a requirement for equi-recursive types. Then, we have the possibility to prove coercions by coinduction, which is crucial in order to equal periods.

We present how to recover the usual rules for reasoning on recursive types [3] by using rule `PROPFIX`. For the usual rule `EQUIVPERIOD`, we use our notion of erasable isomorphism, which we also call equivalence, instead of type equality. This is not a problem because equivalence is interpreted as equality. This rule says that two types τ_1 and τ_2 are equivalent if each of them is equivalent to some unfolding by a common well-founded functor $\alpha \mapsto \sigma$.

$$\frac{\text{EQUIVPERIOD} \quad \alpha \mapsto \sigma : \text{WF} \quad \Sigma; \Theta_0; \Theta_1 \vdash \tau_1 \equiv \sigma[\alpha/\tau_1] \quad \Sigma; \Theta_0; \Theta_1 \vdash \tau_2 \equiv \sigma[\alpha/\tau_2]}{\Sigma; \Theta_0; \Theta_1 \vdash \tau_1 \equiv \tau_2}$$

This rule can be used to show that one-step τ -lists $\mu\alpha 1 + \tau \times \alpha$ are equivalent to two-steps τ -lists $\mu\alpha 1 + \tau \times (1 + \tau \times \alpha)$. It suffices to take τ_1 to be one-step τ -lists, τ_2 to be two-steps τ -lists, and σ to be $1 + \tau \times (1 + \tau \times \alpha)$. To show the first equivalence, we have to unfold τ_1 twice, and to show the second equivalence, we have to unfold τ_2 once.

Rule `EQUIVPERIOD` is derivable for each type σ because the derivation follows the structure of σ . We first use rule `PROPFIX` to memorize the $\tau_1 \equiv \tau_2$ proposition. We then split the conjunction and show only the side $\tau_1 \triangleright \tau_2$, because the other side is similar. We then use both hypotheses to get $\sigma[\alpha/\tau_1] \triangleright \sigma[\alpha/\tau_2]$. We now follow the structure of σ using η -expansions to go under constructors and reflexivity when subtypes are free from the type variable α . When we reach α , it means that we have to show either $\tau_1 \triangleright \tau_2$ or $\tau_2 \triangleright \tau_1$ depending on the variance. But in both cases we went under a computational type constructor, because σ is well-founded with respect to α , and the $\tau_1 \equiv \tau_2$ hypothesis is now accessible. This remark ends the proof.

The second usual rule about recursive types is `COERETAMU`. It tells when two recursive types are in the coercion relation. The recursive type $\mu\alpha \tau$ is smaller than $\mu\beta \sigma$ if τ is smaller than σ given that α is smaller than β . This rule can be used to prove that τ -lists $\mu\alpha 1 + \tau \times \alpha$ are smaller than σ -lists $\mu\beta 1 + \sigma \times \beta$ if τ is smaller than σ . We simply use η -expansions for the sum and product types, and the hypotheses for $\tau \triangleright \sigma$ and $\alpha \triangleright \beta$.

$$\frac{\text{COERETAMU} \quad \Sigma, (\alpha\beta : \star \times \star \mid (\emptyset \vdash \alpha) \triangleright \beta); \Theta_0; \Theta_1 \vdash (\emptyset \vdash \tau) \triangleright \sigma}{\Sigma; \Theta_0; \Theta_1 \vdash (\emptyset \vdash \mu\alpha \tau) \triangleright \mu\beta \sigma}$$

This rule is derivable in System F_{cc} . We use rule `PROPFIX` to memorize the $\mu\alpha \tau \triangleright \mu\beta \sigma$ coercion. We then use fold and unfold to get $\tau[\alpha/\mu\alpha \tau] \triangleright \sigma[\beta/\mu\beta \sigma]$. We now use our body hypotheses and modify its use of $\alpha \triangleright \beta$ by our memorized coercion, which is accessible because $\mu\alpha \tau$ and $\mu\beta \sigma$ are well-formed recursive types and thus $\alpha \mapsto \tau$ and $\beta \mapsto \sigma$ are well-founded.

5.5 Incoherent Polymorphism

Incoherent polymorphism is a necessity for some type system features, but it may also be a simplification. First, as described below in Section 6.1.2, incoherent polymorphism is necessary to describe the kind of existentials used in GADTs (where incoherence is frequent). On the one hand, coherent polymorphism requires the kind to be coherent, and thus the existence of a witness type. On the other hand, incoherent polymorphism permits type abstraction for any well-formed kind: inhabited kinds, potentially inhabited kinds, and empty kinds. In the polymorphic type $\forall(\alpha : \kappa) \tau$, the coherence of kind κ may depend over type variables β of the type environment. Depending on how they are instantiated, the kind κ may or may not be inhabited. A simple example will follow when illustrating GADTs. Apart from these cases,

we may simply omit to prove that the kind κ is coherent to ease the programmers task. The result will be inaccessible code until the abstraction is instantiated. Since such abstraction is equivalent to a unit abstraction (see the last paragraph of this section about weak reduction), we may say it has a zero-bit cost as in FC. Actually, this construct corresponds to the notion of coercion abstraction in FC. In the rest of this section, we describe how to extend System F_{cc} with incoherent polymorphism, and illustrate its use with GADTs. This extension is in the Coq development and thus proved sound.

We extend the syntax of the λ -calculus with type abstraction ∂a and type application $a \diamond$. Type abstraction does not bind anything. Its only use is to block reduction. Hence, we only add the evaluation context for type application $\square \diamond$. The type abstraction context $\partial \square$ is not an evaluation context. This is the most important part of this extension. In particular, ∂r , where r is an error, is sound. Notice that removing an evaluation context may break confluence, which is the case here. But we can restore it by keeping a list of substitutions at type abstractions as it is done in [22] or using an type weakening construct as it is done in [8].

Prevalues and values are extended accordingly: $p \diamond$ is a prevalue and ∂a is a value. We do not ask the body a of the type abstraction to be a value because we do not reduce under type abstraction. Finally, we extend the set of errors with destructors applied to the wrong constructor.

$a, b ::= \dots \mid \partial a \mid a \diamond$	Terms
$E ::= \dots \mid \square \diamond$	Evaluation contexts
$p ::= \dots \mid p \diamond$	Prevalues
$v ::= \dots \mid \partial a$	Values
$r ::= \dots \mid (\partial a) a \mid \text{fst}(\partial a) \mid \text{snd}(\partial a) \mid (\lambda x a) \diamond \mid \langle a, a \rangle \diamond$	Errors

We add one reduction rule for type applications. When a type abstraction occurs right under a type application, both constructs annihilates and unfreeze the body of the type abstraction which can now be reduced.

$$\begin{array}{c} \text{REDTAPP} \\ (\partial a) \diamond \rightsquigarrow a \end{array}$$

We now extend the type system. We extend types with incoherent polymorphism $\Pi(\alpha : \kappa) \tau$. The notation differs from coherent polymorphism $\forall(\alpha : \kappa) \tau$ only by the \forall quantifier which becomes a Π quantifier (that has nothing to do with the quantifier for dependent products).

$$\tau, \sigma ::= \dots \mid \Pi(\alpha : \kappa) \tau \qquad \text{Types}$$

We extend typing rules as follows. We add a rule RECPI to allow recursive type variables to occur in the body of an incoherent polymorphic type. Because incoherent polymorphic types are computational, they are well-founded as long as their body is non-expansive. Notice that we ask the recursive type variable not to be free in the kind of the abstract type variable. Rule TYPEPI is similar to rule TYPEFOR. An incoherent polymorphic type has the star kind if its body has the star kind under an extended environment.

$$\begin{array}{c} \text{RECPI} \\ \frac{\alpha \mapsto \tau : \text{NE} \quad \alpha \notin \text{fv}(\kappa)}{\alpha \mapsto \Pi(\beta : \kappa) \tau : \text{WF}} \end{array} \qquad \begin{array}{c} \text{TYPEPI} \\ \frac{\Sigma, (\alpha : \kappa) \vdash \tau : \star}{\Sigma \vdash \Pi(\alpha : \kappa) \tau : \star} \end{array}$$

Rule `TERMGEN` is the introduction rule for incoherent polymorphism. There are three main differences between `TERMGEN` and `COERGEN`, which is the introduction rule of coherent polymorphism. First, `TERMGEN` is a typing rule for terms, while `COERGEN` is a coercion rule, as their names suggest. This comes from the fact that incoherent polymorphic types are computational, while coherent polymorphic types are erasable. This implies the second difference: `TERMGEN` leaves a mark on the term, while `COERGEN` is erasable. Finally, `TERMGEN` only requires the kind to be well-formed, while `COERGEN` needs the kind to be coherent. These are very important distinctions, because relaxing the coherence condition for rule `COERGEN` would be unsound. Rule `TERMINST` is the elimination rule of incoherent polymorphism. The only difference with rule `COERINST` is erasability: `TERMINST` leaves a mark while `COERINST` is erasable.

$$\frac{\text{TERMG}_{\text{EN}} \quad \Sigma \Vdash \kappa \quad a : \Sigma, (\alpha : \kappa); \Gamma \vdash \tau \quad \alpha \notin \text{fv}(\Gamma)}{\partial a : \Sigma; \Gamma \vdash \Pi(\alpha : \kappa) \tau} \quad \frac{\text{TERMINST} \quad a : \Sigma; \Gamma \vdash \Pi(\alpha : \kappa) \tau \quad \Sigma \vdash \sigma : \kappa}{a \diamond : \Sigma; \Gamma \vdash \tau[\alpha/\sigma]}$$

Since computational types have non-erasable introduction and elimination rules, we cannot derive their η -expansion coercion as we can do with erasable types. Incoherent polymorphic types are computational and we thus need to add an η -expansion coercion rule `COERPI` for them. Like other η -expansion rules, coinductive hypotheses are made accessible. This is due to the computational behavior of incoherent polymorphic types.

$$\frac{\text{COERP}_{\text{I}} \quad \Sigma \Vdash \kappa \quad \Sigma \vdash \Sigma' \quad \alpha \notin \text{fv}(\Theta_0, \Theta_1, \Sigma', \kappa', \tau') \quad \Sigma, (\alpha : \kappa); \Theta_0, \Theta_1; \emptyset \vdash (\Sigma' \vdash \tau'[\alpha'/\sigma']) \triangleright \tau \quad \Sigma, (\alpha : \kappa), \Sigma' \vdash \sigma' : \kappa'}{\Sigma; \Theta_0; \Theta_1 \vdash (\Sigma' \vdash \Pi(\alpha' : \kappa') \tau') \triangleright \Pi(\alpha : \kappa) \tau}$$

To understand the main ideas of this somewhat complicated rule, we have to draw its explicit coercion as an η -expansion context. We use the line between nodes for the turnstile: type environments are on the left and types are on the right. We start to read this typing derivation from the bottom, which is the resulting typing of the η -expansion $\Sigma \vdash \Pi(\alpha : \kappa) \tau$. We first use `TERMGEN` which moves the type binding from the type to the environment to get the typing $\Sigma, (\alpha : \kappa) \vdash \tau$. We then use the subcoercion and assume it binds Σ' . We do not do any assumption about the argument type of the subcoercion to be principal. We find it by using `TERMINST` which gives a type substitution of α' in τ' for some σ' . By typing, we deduce that we need $\Sigma, (\alpha : \kappa), \Sigma' \vdash \sigma' : \kappa'$. Finally, we realize that we need to use a weakening rule to remove the extra binding $(\alpha : \kappa)$ to get the argument typing of the η -expansion $\Sigma, \Sigma' \vdash \Pi(\alpha' : \kappa') \tau'$. The weakening rule requires α not to be free in Σ', κ' , and τ' (plus the coinduction hypotheses Θ_0 and Θ_1 which are not drawn).

$$\begin{array}{c}
\Sigma, \Sigma' \downarrow \Pi(\alpha' : \kappa') \tau' \\
\text{weak} \\
\Sigma, (\alpha : \kappa), \Sigma' \downarrow \Pi(\alpha' : \kappa') \tau' \\
\circlearrowleft \leftarrow \Sigma, (\alpha : \kappa), \Sigma' \vdash \sigma' : \kappa' \\
\Sigma, (\alpha : \kappa), \Sigma' \downarrow \tau'[\alpha'/\sigma'] \\
\text{coer} \\
\Sigma, (\alpha : \kappa) \downarrow \tau \\
\Lambda \\
\Sigma \downarrow \Pi(\alpha : \kappa) \tau
\end{array}$$

The well-formedness of kind κ is for extraction purposes. Notice that we do not require coherence for the kind κ because this is the η -expansion of the incoherent polymorphic type. However, we require the coherence of the type environment extension Σ' under Σ . This is a very important premise because we do not want the incoherence of κ to leak in Σ' and thus under the coercion, because coercions are erasable.

A simple counter-example would be to use rule **COERGEN** to prove $\Sigma, (\alpha : \kappa); \emptyset; \emptyset \vdash (\emptyset, (\beta : \kappa) \vdash \tau) \triangleright \forall(\beta : \kappa) \tau$, which holds using α as a witness. We could then give this sub-coercion to rule **COERP1** with instantiation type α to get $\Sigma; \emptyset; \emptyset \vdash (\emptyset, (\beta : \kappa) \vdash \Pi(\alpha : \kappa) \tau) \triangleright \Pi(\alpha : \kappa) \forall(\beta : \kappa) \tau$. We could thus build an erasable coercion that extends the environment with a potentially incoherent binding like $(\beta : 1 \mid (\emptyset \vdash \mathbf{Bool}) \triangleright \mathbf{Int})$.

The indexed calculus is extended accordingly to the extensions of the λ -calculus. We define the interpretation of $\Pi(\alpha : \kappa) \tau$ under η as a variant of the arrow type and the coherent polymorphic type, because it is a computational type and it behaves as an intersection in terms of typing. The main difference is that we do not ask the inner term to be sound. All these extensions are formalized in **Coq** and thus proved sound and normalizing (Theorem 101 and 102).

$$|\Pi(\alpha : \kappa) \tau|_{\eta} \stackrel{\text{def}}{=} \diamond \left\{ \partial^k e \mid k > 0 \Rightarrow \forall x \in |\kappa|_{\eta} \ [e]_{k-1} \in |\tau|_{\eta, \alpha \mapsto x} \right\}$$

Let's illustrate a practical case when incoherent polymorphism is useful. Let's first define existentials by CPS encoding (see Section 6.1.2). Because we have two notions of polymorphism, coherent and incoherent, we also have two notions of existentials: coherent and incoherent. We write $\exists(\alpha : \kappa) \tau$ for coherent existential types and $\Sigma(\alpha : \kappa) \tau$ for incoherent existential types with the following definitions (using the CPS encodings):

$$\begin{array}{ll}
\text{coherent:} & \exists(\alpha : \kappa) \tau \stackrel{\text{def}}{=} \forall \beta (\forall(\alpha : \kappa) (\tau \rightarrow \beta)) \rightarrow \beta \\
\text{incoherent:} & \Sigma(\alpha : \kappa) \tau \stackrel{\text{def}}{=} \forall \beta (\Pi(\alpha : \kappa) (\tau \rightarrow \beta)) \rightarrow \beta
\end{array}$$

We define the **pack** and **unpack** term syntactic sugar for the coherent existential, and **ipack** and **iunpack** for their incoherent version. Notice that the body of the **iunpack** sugar is hidden under an incoherent type abstraction, and as such is allowed to be unsound because it cannot be reduced.

$$\begin{array}{lll}
\text{coherent:} & \mathbf{pack} \ a \stackrel{\text{def}}{=} \lambda x \ x \ a & \mathbf{unpack} \ a \ \text{as} \ x \ \text{in} \ b \stackrel{\text{def}}{=} a \ (\lambda x \ b) \\
\text{incoherent:} & \mathbf{ipack} \ a \stackrel{\text{def}}{=} \lambda x \ x \ \diamond a & \mathbf{iunpack} \ a \ \text{as} \ x \ \text{in} \ b \stackrel{\text{def}}{=} a \ (\partial \lambda x \ b)
\end{array}$$

Let's assume we have type-level functions (see Section 6.1.3) and sum types as in Section 2.4.3. We also use the erasable isomorphism notation of Section 5.4.1. We can now define the following GADT, named **Term**, and with kind $\star \rightarrow \star$. We first give its Haskell syntax:


```

data Term a where
  Lam :: (a -> b) -> Term (a -> b)
  App :: Term (a -> b) -> Term a -> Term b

```

In our mathematical syntax, we write:

$$\text{Term } \alpha \stackrel{\text{def}}{=} \Sigma(\beta : \star \times \star \mid \alpha \equiv (\text{fst } \beta \rightarrow \text{snd } \beta)) \alpha \\ + \exists \beta \text{ Term } (\beta \rightarrow \alpha) \times \text{Term } \beta$$

This GADT is the sum of an incoherent existential type and a coherent existential type. The incoherent existential type asks α to be an arrow type and stores a term of such type. It also names $\text{fst } \beta$ the argument type and $\text{snd } \beta$ its return type. The coherent existential type adds no constraint on α but stores a pair such that its first component applied to its second component is of type α . It names β the intermediate type. The **Term** GADT contains two constructors: one for the left-hand side of the sum injecting functions and one for the right-hand side of the sum freezing function applications. We can define its two constructors in the following manner:

$$\text{Lam } x \stackrel{\text{def}}{=} \text{inl } (\text{ipack } x) \quad : \forall \alpha \forall \beta (\alpha \rightarrow \beta) \rightarrow \text{Term } (\alpha \rightarrow \beta) \\ \text{App } y \ x \stackrel{\text{def}}{=} \text{inr } (\text{pack } \langle y, x \rangle) : \forall \alpha \forall \beta \text{ Term } (\alpha \rightarrow \beta) \rightarrow \text{Term } \alpha \rightarrow \text{Term } \beta$$

We can now define a recursive eval function taking a term of type **Term** α and returning a term of type α for all type variable α . Said otherwise, **eval** has type $\forall \alpha \text{ Term } \alpha \rightarrow \alpha$. When the argument is on the left-hand side of the sum, **eval** simply unpacks the inside argument and returns it. When the argument is on the right-hand side of the sum, **eval** unpacks the inside argument, which is a pair. It applies the evaluation of the first component to the evaluation of the second component. On the left-hand side, we use the incoherent version of unpack, while we use the coherent version on the right-hand side.

```

eval :: Term a -> a
eval arg = case arg of
  Lam f -> f -- this branch is potentially incoherent
  App f x -> (eval f) (eval x)

```

$$\text{eval } x = \text{case } x \text{ of } \{ \text{inl } x_1 \mapsto \text{iunpack } x_1 \text{ as } y \text{ in } y \\ | \text{inr } x_2 \mapsto \text{unpack } x_2 \text{ as } y \text{ in } (\text{eval } (\text{fst } y)) (\text{eval } (\text{snd } y)) \}$$

Let's now suppose that we call **eval** with a term of type **Term** $(\tau \times \sigma)$. This term is necessarily from the right-hand side of the sum because $\tau \times \sigma$ cannot be equivalent to an arrow type by consistency. However, in the first branch, in the body of the inconsistent unpack, we have access to the proposition $\tau \times \sigma \equiv \text{fst } \beta \rightarrow \text{snd } \beta$ which is inconsistent. This sort of inconsistency in some branches of case expression on GADTs is frequent. Notice however, that in the second branch we can reduce for any instantiation of α because we used a coherent existential type: there is a witness for β for any instance of α .

Weak reduction If the language is equipped (in addition to our current strong term abstraction) with a weak term abstraction, *i.e.* a term abstraction under which reduction is not allowed, then it is possible to reuse this existing abstraction to implement incoherent type

abstraction. Let's write $\Lambda x a$ this weak term abstraction. It can use the existing application construct, because weak reduction is a property of the abstraction, not the application. We do not extend evaluation contexts, and in particular we do not add $\Lambda x []$. We add the value $\Lambda x a$ for weak term abstractions. And we extend errors accordingly. We extend the type system with the weak arrow type $\Pi(\alpha : \kappa)\tau \Rightarrow \tau$, which does a type abstraction first, and then a weak term abstraction.

$a, b ::= \dots \mid \Lambda x a$	Terms
$v ::= \dots \mid \Lambda x a$	Values
$r ::= \dots \mid \mathbf{fst}(\Lambda x a) \mid \mathbf{snd}(\Lambda x a)$	Errors
$\tau, \sigma ::= \dots \mid \Pi(\alpha : \kappa)\tau \Rightarrow \tau$	Types

We add a reduction rule similar to REDAPP. The main difference is in the typing rule of this term abstraction. We also need to add a typing rule for application of such abstraction. Rule TERMLAMWEAK tells that $\Lambda x a$ has type $\Pi(\alpha : \kappa)\tau \Rightarrow \sigma$ under Σ and Γ , if α is not free in Γ and if the body a has type σ after type abstraction $(\alpha : \kappa)$ and term abstraction $(x : \tau)$. Rule TERMAPPWEAK tells that $a b$ has type $\tau_2[\alpha/\sigma]$, if a has type $\Pi(\alpha : \kappa)\tau_1 \Rightarrow \tau_2$, σ has kind κ , and b has type $\tau_1[\alpha/\sigma]$.

$$\begin{array}{c}
 \text{REDAPPWEAK} \\
 (\Lambda x a) b \rightsquigarrow a[x/b]
 \end{array}
 \quad
 \frac{
 \begin{array}{c}
 \text{TERMLAMWEAK} \\
 \Sigma \Vdash \kappa \quad \Sigma, (\alpha : \kappa) \vdash \tau : \star \quad a : \Sigma, (\alpha : \kappa); \Gamma, (x : \tau) \vdash \sigma \quad \alpha \notin \text{fv}(\Gamma)
 \end{array}
 }{
 \Lambda x a : \Sigma; \Gamma \vdash \Pi(\alpha : \kappa)\tau \Rightarrow \sigma
 }$$

$$\frac{
 \begin{array}{c}
 \text{TERMAPPWEAK} \\
 a : \Sigma; \Gamma \vdash \Pi(\alpha : \kappa)\tau_1 \Rightarrow \tau_2 \quad \Sigma \vdash \sigma : \kappa \quad b : \Sigma; \Gamma \vdash \tau_1[\alpha/\sigma]
 \end{array}
 }{
 a b : \Sigma; \Gamma \vdash \tau_2[\alpha/\sigma]
 }$$

We can now simulate type abstraction $\Pi(\alpha : \kappa)\tau$ with $\Pi(\alpha : \kappa)1 \Rightarrow \tau$ and we can also simulate the usual weak term abstraction $\tau \Rightarrow \sigma$ with $\Pi(\alpha : 1)\tau \Rightarrow \sigma$. This enhances the fact that incoherent type abstractions are actually unit weak abstractions. Both views are actually equivalent for well-typed terms. We can encode the weak abstraction $\Lambda x a$ using type abstraction and strong abstraction to form $\partial \lambda x a$. And we can encode the application $a b$ for weak abstraction (which we know by typing) using type application and application to form $a \diamond b$. Reciprocally, we can encode the type abstraction ∂a as $\Lambda x a$ for some fresh term variable x . And we can encode the type application $a \diamond$ as $a \langle \rangle$, where $\langle \rangle$ is the unit term.

5.6 Coq formalization

The Coq [11] formalization has been developed with Coq version 8.4pl2. It can be found online at <http://phd.ia0.fr/>. There is a Makefile, so it suffices to run `make` to compile the Coq files or `make html` to also build the html version. The formalization merges ideas from strong normalization proofs of System F and step-indexed techniques. The strong normalization proof techniques is initially inspired from a strong normalization proof of System F in ATS [14] but adapted for soundness proofs and step-indices. The step-indexed techniques are inspired from a soundness proof with unrestricted recursive types and side effects [4] but modified for strong reduction.

The main differences between the version of System F_{cc} presented in this chapter and its formalization in Coq are the use of de Bruijn indices and the parametrization of the type

system. Using de Bruijn indices makes it a lot easier and cleaner to deal with binders. The parametrization is necessary to state the results and for the induction to go well.

We will now give a small glimpse of the Coq code for the reader to find its way through the files, definitions, and lemma. Files prefixed with the capital letter **F** refer to the indexed calculus: this letter stands for fuel. Files prefixed with the capital letter **L** refer to the lambda calculus. Files without a prefix letter are more general, like `typesystem.v` which factorizes the type systems for both the indexed calculus and the lambda calculus. We describe the files in dependency order.

We first have a few independent files. The file `ext.v` defines the extensionality axioms we use. Propositional and functional extensionality are the only axioms used. The other extensionality rules are lemmas. The file `set.v` defines a type for potentially infinite sets as predicates in `Prop`. The file `minmax.v` defines the tactic `minmax` to deal with indices. The file `list.v` defines useful lemmas about lists, which we use for environments and mappings.

Finally, the file `mxx.v` defines the parametrization of the type system using a value of type `Mode`, which is the pair of a boolean and a version. The boolean tells whether we allow recursive types or not. The soundness proof does not constrain this boolean while the normalization proof requires the boolean to be false. The version is defined by the inductive `Version` and contains three possible values: `vP`, `vF`, and `vS`. This manuscript corresponds to the `vP` version, which is the natural presentation. The `vF` version contains additional premises that are redundant by extraction but necessary for the soundness induction. Finally, the `vS` version removes some premises from the `vF` version which are required for extraction but not for soundness. The proof of soundness is thus done in the `vS` version. We define two helpers to tell whether a premise is required for extraction with `mE` or for soundness with `mS`. Finally, the `mR` helper tells which rules are about recursive types.

We can now define the indexed calculus in file `Flanguage.v`. We define the inductive `term` for terms. All constructors of this inductive are prefixed with a `nat` representing the index of the node. We then define a few functions to traverse terms, from which we define lifting of index predicates to term, and the `lift` and `subst` functions for de Bruijn lifting and substitution. We then define the reduction relation in the inductive `red`. We finally define errors in `Err` and valid terms `V`. What follows is a list of lemmas about lifting, substitutions, lowering, and other functions over indices.

We prove the strong normalization of the pure indexed calculus in file `Fnormalization.v`. This file is quite simple to follow: we define a `measure`, prove that it strictly decreases with reduction in `red_measure`, and finally prove that reduction is well founded in `wf_der`.

We can now define a semantics for this indexed calculus in file `Fsemantics.v`. We define the notion of interior in `Dec`, the notion of contraction in `Red`, and the notion of expansion in `Exp`. Using expansion, we can define the set of sound terms `OK`. We define pretypes in `C` and types in `CE`. We define the closure of a set in `C1`, in order to define the arrow operator `EArr`, product operator `EProd`, and incoherent abstraction operator `EPi`. We show that these operators preserve types in `CE_EArr`, `CE_EProd`, and `CE_EPi`. We also define erasable types such as the coherent polymorphic type `EFor`, the top type `ETop`, the bottom type `EBot`, or recursive types `EMu`. We then define the notions we need to show that recursive types are equal to their unfolding. And we finally define the semantic judgment `EJudg` and the semantic typing rules of the STLC, such as `ELam_sem`. We also define a subtyping rule `ECoer_sem` and a distributivity rule `Edistrib` which will be used together to prove rule `TERMCOER`.

Once that the indexed calculus is defined, we may define the lambda calculus and the functions to go back and forth between them in file `Llanguage.v`. The structure of this

file is similar to `Flanguage.v` with the difference that it now contains a `drop` and `kfill` function to translate terms from one language to the other. It also contains the key lemma `drop_red_exists` for the bisimulation between the reduction relation of both languages.

Independently from the indexed calculus and the lambda calculus, we can define the type part of System F_{cc} : everything but the term judgment. This is done in `typesystem.v` and is actually shared by both `Ftypesystem.v` and `Ltypesystem.v`, which define the term judgment for the indexed type system and lambda type system respectively. The last two files are exactly the same up to indices. All the syntax is gathered in a single Coq inductive, namely `obj`. This simplifies a lot the treatment of operations on the syntax, such as lifting or substitution, which are defined only once. In order to keep track of syntactical classes, we define a judgment `cobj` to classify each object in its grammatical class. In the paper version, we naturally assume that everything is syntactically well-formed, while we have to state it explicitly in Coq.

We prove the weakening, substitution, and extraction lemmas in `typesystemextra.v`. This file also contains the proof that the `vP` and `vF` version are equivalent and that the `vS` version is a consequence. This explains why the properties of the `vS` version also hold for the `vP` version.

The proof that each judgment of the indexed type system is sound lies in `Fsoundness.v`. We start by defining a notion of semantic objects `sobj`. A semantic object is either a set of indexed terms, the unit object, or a pair of objects. We then define the signature of the interpretation of each syntactical class in `sem`. Kinds are interpreted as sets of semantic objects, types as semantic objects, propositions as indexed propositions, type environments as sets of semantic environments (lists of semantic objects because we use de Bruijn indices), coinduction environments as indexed propositions, and finally term environments as semantic term environments (lists of semantic types, when the syntactical environment is valid). All these interpretation are parametrized by an surrounding semantic environment. We define the interpretation function `Fsoundness.semobj` as a binary relation, but we show in `semobj_eq` that it behaves as a function. We then prove semantic lifting and substitution properties. And we finally prove the soundness of each judgment. The soundness of `jfoo` is proved in `jfoo_sound` (Lemma 100).

We finally lift this soundness proof from the indexed type system to the lambda type system in `Lsoundness.v`. We first define when a term a is sound for a least k steps in `OKstep` and when it is sound for all number of steps in `OK` by coinduction. We prove that these two notions coincide. We then show how to transpose soundness from the indexed calculus to the lambda calculus in `term_ge_OK` (Lemma 76). We prove that both type systems coincide and finally prove the soundness of System F_{cc} in `soundness` (Theorem 101).

The Coq files `Lsemantics.v` and `Lnormalization.v` are similar to the files `Fsemantics.v`, `Fsoundness.v`, and `Lsoundness.v` but deal with the strong normalization result instead of the soundness result.

5.7 Discussion about the explicit version

What System F_{cc} lacks in order to have an explicit version is subject reduction. A typical example where we lose subject reduction is inspired from Section 4.6.2 and involves a coercion variable in-between a redex, which requires decomposing coercions to remain well-typed after reduction. Let's take the following definitions:

- $\Sigma \stackrel{\text{def}}{=} \emptyset, (\alpha\beta \mid (\emptyset \vdash \text{Int} \rightarrow \alpha) \triangleright \text{Int} \rightarrow \beta)$

- $\Gamma \stackrel{\text{def}}{=} \emptyset, (y : \alpha)$
- $a \stackrel{\text{def}}{=} (\lambda x y) 3$

We first need to show that Σ is coherent. To do so, it suffices to take \perp as a witness for α and \top as a witness for β . We have $a : \Sigma; \Gamma \vdash \beta$ and $a \rightsquigarrow y$, but we do not have $y : \Sigma; \Gamma \vdash \beta$. We would have to prove $\alpha \triangleright \beta$ from our hypothesis $\text{Int} \rightarrow \alpha \triangleright \text{Int} \rightarrow \beta$. This decomposition on the right-hand side of the arrow type is actually semantically valid, and it is possible to add it and thus restore subject reduction for this example. However, it is unclear whether the decomposition on the left-hand side holds (see later in this section). We can adapt the counter-example to subject reduction to use the left-hand side decomposition. We take the following definitions:

- $\Sigma \stackrel{\text{def}}{=} \emptyset, (\alpha\beta \mid (\emptyset \vdash \alpha \rightarrow \alpha) \triangleright \beta \rightarrow \alpha)$
- $\Gamma \stackrel{\text{def}}{=} \emptyset, (y : \beta)$
- $a \stackrel{\text{def}}{=} (\lambda x x) y$

The type environment Σ is coherent by taking \top as the witness of α and \perp as the witness for β . We have $a : \Sigma; \Gamma \vdash \alpha$ and $a \rightsquigarrow y$, but we do not have $y : \Sigma; \Gamma \vdash \alpha$. We would have to prove $\beta \triangleright \alpha$ from our hypothesis $\alpha \rightarrow \alpha \triangleright \beta \rightarrow \alpha$.

We can prove the decomposition for the right-hand side of the arrow type when the coercion is not binding anything. This means that the initial coercion is simply of the form $\tau' \rightarrow \sigma' \triangleright \tau \rightarrow \sigma$ and not of the form $(\Sigma' \vdash \tau' \rightarrow \sigma') \triangleright \tau \rightarrow \sigma$. From this coercion we can extract a coercion from σ' to σ . If the coercion was binding something, using a distributivity rule prior to this lemma should work.

Lemma 104 (Push right). *If S' and R are types and $R' \rightarrow S' \subseteq R \rightarrow S$ holds, then $S' \subseteq S$ holds.*

Proof. Coq lemma `Fsemantics.Push_right_Arr`

Let $e \in S'$ and show that $e \in S$. Let k be the greatest index in e . We have $e \leq k$. Let x be a fresh variable, and in particular not free in e . We have $\lambda^{k+1} x e \in R' \rightarrow S'$ by definition of the arrow operator and because $[e]_k = e$. By hypothesis, we also have $\lambda^{k+1} x e \in R \rightarrow S$. And because all semantic types are inhabited, we use the arrow operator definition to get $[e]_k \in S$ which concludes. \square

It is less clear whether or not the decomposition on the left-hand side of the arrow operator holds. However, we can show that this decomposition does not hold in other semantics, like in weak reduction for example. Some semantics only consider closed terms and some others permit semantic types to be empty. One example of a semantics satisfying these two assertions is when semantic types are sets of closed strong normalizing terms. In this setting, the bottom type is the empty set. And functions types from inhabited types to the bottom type are empty too. In particular $\text{Int} \rightarrow \perp$ is empty and thus included in $\text{Bool} \rightarrow \perp$ which is also empty. If we could decompose coercions on the left of arrow types, we would get $\text{Bool} \triangleright \text{Int}$ which is unsound.

So, for some settings, coercion decomposition is unsound. For our setting, this is an open question, although we have a positive result for the right-hand side of the arrow operator.

And finally, for more restrictive settings, coercion decomposition is sound and thus part of the system. Such restrictive setting is **FC** that considers only equality coercions. Equality coercions enforce structural coercions which simplifies the framework, as the main difficulty comes from type generalization and instantiation being coercions. Another less restrictive simplification would be to remove polymorphism from the coercion judgment and use it only in the typing judgment. However, we would lose a lot of the current expressivity and generality.

Chapter 6

Discussions

This work should be applicable to other programming languages instead of the λ -calculus, because the idea is to say that once we defined approximations (or invariants) for programs, it is natural to study the order relation between these approximations, define a syntactical composable judgment to prove inclusions between invariants, and study how to abstract over this judgment. However what can be done in the λ -calculus has not yet been fully explored. Several extensions remain to be studied: from new erasable types to changes of the framework. These extensions are described in Section 6.1. We then discuss related works in Section 6.2, and applications in Section 6.3.

6.1 Extensions

In this section, we discuss a series of extensions. Some of them are only conjectures, while some have more technical materials. These materials may either be ideas to further study or strong candidates waiting to be proved. We detail for each extension in which state it lies and which parts may be not quite correct as stated.

6.1.1 Data types

The λ -calculus can be extended with sums $\tau + \sigma$, integers `Int`, booleans `Bool`, the unit type `1`, and the void type `0`. Extending System F_t^p and System F_{cc} to handle `Int`, `Bool`, and `1` is anecdotal. The Coq development of System F_{cc} already enjoys sum types, the unit type, and the void type.

6.1.2 Existentials

There are two ways to add existentials: one is to use their CPS encoding with polymorphic types and another one is to give them their own meaning as union. The first solution always work while the second one may raise difficulties.

Existentials as CPS encoding The well-known encoding from existentials with polymorphism also works in System F_{cc} . We can define the type $\exists(\alpha : \kappa) \tau$ and add the following type equality (or notation): $\exists(\alpha : \kappa) \tau = \forall \beta (\forall (\alpha : \kappa) \tau \rightarrow \beta) \rightarrow \beta$. We can also define `pack a` as a sugar for $\lambda x x a$ and `unpack a as x in b` as a sugar for $a (\lambda x b)$.

This easy extension is actually enough to express GADTs using incoherent polymorphism like in FC. See Section 5.5 for an example of GADTs using incoherent polymorphism, this encoding of existentials, and sum types.

Existentials as unions By contrast with intersection of types, union of types is not obviously a type in our semantics. Our semantics is the usual semantics for reducibility candidates modified with step-indices. Without the expansion-closure, it is not obvious that the union $R \cup S$ of two types is a type. The problem is that a term e in $\diamond(R \cup S)$ could a priori reduce to both e_1 in $\Delta R \setminus \Delta S$ and e_2 in $\Delta S \setminus \Delta R$ and then not be in $R \cup S$.

In the current setting, where the underlying programming language is the λ -calculus, it seems that e should always be in $R \cup S$ by a standardization argument [33]. However, this argument is already complex in the absence of indices and may not be applicable in the case of indices—or force us to have a more involved definition of indexing compatible with standardization. This is an interesting question to explore further.

An alternative to the standardization argument, which would also restore subject reduction, is to use a call-by-constructor strategy allowing destructors to go under redexes. This reduction is equivalent to strong reduction but only substitutes constructors and can be inspired from [2]. Substituting only constructors permit to know that the term is in the kernel of the semantics and not its closure, which is necessary to know in which side of the union a term lies.

6.1.3 Type-level functions

Type-level functions and recursive types at arbitrary kinds (Section 6.1.4) are the last element we need to feature user data-types. User defined data-types are usually a recursive type of a series of type abstractions followed by a sum of products. For instance, the **Term** GADT defined in Section 5.5 does a recursive call applied to an arrow type in the right side of the sum. It should be possible to extend System F_{cc} with type functions, but nothing has been proved yet.

System F_ω (Chapter 30 of [29]) is the simplest type system featuring polymorphism and type-level functions. In order to extend System F_{cc} with type-level functions, we need to extend syntactical types with the type abstraction $\lambda(\alpha : \kappa) \tau$ and type application $\tau \sigma$. We also need to extend syntactical kinds with the arrow kind $\kappa \rightarrow \kappa$. We add the β -reduction rule at the type-level as an equality between types: $(\lambda(\alpha : \kappa) \tau) \sigma = \tau[\alpha/\sigma]$. We finally give the two kinding rules for type abstraction and application, and the well-formedness rule for the arrow kind.

$$\begin{array}{c}
\text{TYPELAM} \\
\frac{\Sigma \Vdash \kappa_1 \quad \Sigma, (\alpha : \kappa_1) \vdash \tau : \kappa_2}{\Sigma \vdash \lambda(\alpha : \kappa_1) \tau : \kappa_1 \rightarrow \kappa_2}
\end{array}
\qquad
\begin{array}{c}
\text{TYPEAPP} \\
\frac{\Sigma \vdash \tau : \kappa_1 \rightarrow \kappa_2 \quad \Sigma \vdash \sigma : \kappa_1}{\Sigma \vdash \tau \sigma : \kappa_2}
\end{array}
\qquad
\begin{array}{c}
\text{WFKARR} \\
\frac{\Sigma \Vdash \kappa_1 \quad \Sigma \Vdash \kappa_2}{\Sigma \Vdash \kappa_1 \rightarrow \kappa_2}
\end{array}$$

This is all we need in order to add type-level functions. We can recover subtyping rules and the notion of variance with the notions we already have: polymorphic propositions and coercions. We do so by lifting our notion of coercions to higher kinds, as we did with isomorphisms in Section 5.4.1. However, the definition of subtyping at the arrow kind is not canonical and depends on the variance we want to consider. So we will index our subtyping relation with a vkind (variance kind), which is a kind where arrow kinds are annotated with

a variance. Type systems with polarized higher-order subtyping [34] have a common notion for kinds and v kinds. We believe these are two distinct notions and keep them separated.

There are at least four possible variances, written d : $+$ for covariant, $-$ for contravariance, $=$ for invariance, and \emptyset for bivarience. V kinds, written K , are the star v kind \star , the unit v kind 1 , the product v kind $K \times K$, the constrained v kind $\{\alpha : K \mid P\}$, and the arrow v kind $K[\mathsf{d}] \rightarrow K$. We write $\lfloor K \rfloor$ the function from v kinds to kinds that drops the variance annotations. The only interesting case is $\lfloor K_1[\mathsf{d}] \rightarrow K_2 \rfloor$ which is defined by $\lfloor K_1 \rfloor \rightarrow \lfloor K_2 \rfloor$.

We write $\tau \triangleright \sigma : K$ to say that τ is smaller than σ at v kind K . Both types τ and σ have kind $\lfloor K \rfloor$. We may consider heterogeneous subtyping later, and focus on homogeneous subtyping for simplicity. Notice that it makes no sense to talk about binding subtyping (or coercions) at higher kinds, because the notion of binding is a notion of terms and thus only relevant at the kind star. The subtyping at v kind K is inductively defined. The first cases are straightforward:

$$\begin{aligned} \tau \triangleright \sigma : \star & \stackrel{\text{def}}{=} (\emptyset \vdash \tau) \triangleright \sigma \\ \tau \triangleright \sigma : 1 & \stackrel{\text{def}}{=} \top \\ \tau \triangleright \sigma : K_1 \times K_2 & \stackrel{\text{def}}{=} (\text{fst } \tau \triangleright \text{fst } \sigma : K_1) \wedge (\text{snd } \tau \triangleright \text{snd } \sigma : K_2) \\ \tau \triangleright \sigma : \{\alpha : K \mid P\} & \stackrel{\text{def}}{=} \tau \triangleright \sigma : K \end{aligned}$$

For the arrow v kind, there is one definition for each possible variance:

$$\begin{aligned} \tau \triangleright \sigma : K_1[+] \rightarrow K_2 & \stackrel{\text{def}}{=} \forall(\alpha\beta : \lfloor K_1 \rfloor \times \lfloor K_1 \rfloor \mid \alpha \triangleright \beta : K_1) \tau \alpha \triangleright \sigma \beta : K_2 \\ \tau \triangleright \sigma : K_1[-] \rightarrow K_2 & \stackrel{\text{def}}{=} \forall(\alpha\beta : \lfloor K_1 \rfloor \times \lfloor K_1 \rfloor \mid \beta \triangleright \alpha : K_1) \tau \alpha \triangleright \sigma \beta : K_2 \\ \tau \triangleright \sigma : K_1[=] \rightarrow K_2 & \stackrel{\text{def}}{=} \forall(\alpha\beta : \lfloor K_1 \rfloor \times \lfloor K_1 \rfloor \mid \alpha \triangleright \beta : K_1 \wedge \beta \triangleright \alpha : K_1) \tau \alpha \triangleright \sigma \beta : K_2 \\ \tau \triangleright \sigma : K_1[\emptyset] \rightarrow K_2 & \stackrel{\text{def}}{=} \forall(\alpha\beta : \lfloor K_1 \rfloor \times \lfloor K_1 \rfloor) \tau \alpha \triangleright \sigma \beta : K_2 \end{aligned}$$

We recover the usual definition that a type τ of kind $\star \rightarrow \star$ has variance d if the subtyping proposition $\tau \triangleright \tau : \star[\mathsf{d}] \rightarrow \star$ holds, by looking at the diagonal of the subtyping relation. We can actually see each v kind K as a constrained kind $\{\alpha : \lfloor K \rfloor \mid \alpha \triangleright \alpha : K\}$ as it is done in polarized higher-order subtyping systems. This definition permits to talk about the kind $\star[+] \rightarrow \star$ of covariant type functions, for instance.

We also recover the well-known facts that a bivariate type is also covariant and contravariant, and that covariant or contravariant types are also invariant. We can actually show in System F_{cc} the proposition that bivarience implies covariance (and also the other implications), which gives us access to variance demoting. The proposition is:

$$\forall(\alpha : 1 \mid \tau \triangleright \sigma : K_1[\emptyset] \rightarrow K_2) \tau \triangleright \sigma : K_1[+] \rightarrow K_2$$

This is a polymorphic proposition over the constrained kind that contains the unit type $\langle \rangle$ of kind 1 if τ is in the bivarience relation with σ at kind $K_1 \rightarrow K_2$. The proof uses rule `PROPGEN` to introduce both $(\alpha : 1 \mid \tau \triangleright \sigma : K_1[\emptyset] \rightarrow K_2)$ and $(\alpha\beta : \lfloor K_1 \rfloor \times \lfloor K_1 \rfloor \mid \alpha \triangleright \beta : K_1)$ in the type environment. We show $\tau \alpha \triangleright \sigma \beta : K_2$ by first fetching α with rule `TYPEVAR`, then applying rule `PROPRES` to extract its proposition $\forall(\alpha\beta : \lfloor K_1 \rfloor \times \lfloor K_1 \rfloor) \tau \alpha \triangleright \sigma \beta : K_2$, and finally using `PROPINST` with α and β .

6.1.4 Recursive types at arbitrary kinds

Our current version of recursive types is of the form $\mu\alpha \tau$ and only permits to build types of the star kind. Moreover, the folding and unfolding of recursive types are coercions and not type

equalities. We showed in Section 5.4.4 that we do not lose the power of equi-recursive types by using coercions. But the kind star restriction for recursive types is a real limitation when the type system has other sorts of kinds. Recursive types at product kinds permit to build mutually recursive types, while arrow kinds permit to build data-structures (see Section 6.1.3). We do not know how to exactly extend System F_{cc} with recursive types of arbitrary kinds, but we have a privileged path to follow, which we describe below.

Recursive types at arbitrary kind would be written $\mu(\alpha : \kappa)\tau$. Since their unfolding rule is at kind κ , it cannot be a coercion anymore and has to be defined as a type equality $\mu(\alpha : \kappa)\tau = \tau[\alpha/\mu(\alpha : \kappa)\tau]$. This implies that the type equality judgment needs to check the well-foundedness of recursive types, for the folding and unfolding to be sound. The kinding rule for the recursive type has to be modified to take the kind into account:

$$\frac{\text{TYPEMU} \quad \alpha \mapsto \tau : \mathbf{WF} \quad \Sigma, (\alpha : \kappa) \vdash \tau : \kappa}{\Sigma \vdash \mu(\alpha : \kappa)\tau : \kappa}$$

We also need to extend the notion of well-founded and non-expansive functors to all type constructs. This extension will follow the intuition that all occurrences of a recursive type variable have to cross a computational type for the recursive type to be well-formed. To prove the soundness of such rules, we extend the notion of k -approximations to all mathematical objects: not only sets of terms in System F_{cc} , but also to the unit object, pairs of objects, and functions on objects (see Section 6.1.3). We would also need to define the limit of a series of objects when k approaches infinity.

A notion of sort has to be added for mathematical objects: sets, unit, products, and arrows. This corresponds to a simply-typed object calculus. The k -approximation and limit operators would depend on the sort of objects we give them. We have four k -approximation:

- $\langle R \rangle_k^{\text{set}} = \{e \in R \mid e < k\}$
- $\langle R \rangle_k^1 = \langle \rangle$
- $\langle R \rangle_k^{s_1 \times s_2} = \langle \langle \text{fst } R \rangle_k^{s_1}, \langle \text{snd } R \rangle_k^{s_2} \rangle$
- $\langle R \rangle_k^{s_1 \rightarrow s_2} = X \mapsto \langle R X \rangle_k^{s_2}$ or $X \mapsto \langle R \langle X \rangle_k^{s_1} \rangle_k^{s_2}$

We η -expand the objects according to their sort and use the approximation for their subsorts for their subcomponents. We use the same mechanism for limits. We have not fully explored this idea and do not know whether or not the details can be worked out. It may also require that type functions be non-expansive.

Finally, it would be interesting to add the well-foundedness judgment as a proposition assertion and thus internalize well-foundedness hypotheses. This would permit to write an abstract module for lists, sets, or mappings and still remember that these functors are well-founded.

6.1.5 Non-erasable coercions

In this thesis, we have only studied erasable coercions. What would be a more general notion of coercions? There are several ways to extend erasable coercions to non-erasable ones and generalize the notion of coercions. We did not study non-erasable coercions, although we feature incoherent polymorphism.

A first example of non-erasable coercions is given in FC. Instead of being fully erased, coercions are erased to the unit term. Abstracting over a coercion results in a simple abstraction without bindings because the unit element contains no information. The information of a coercion is only present at typing. This view is compatible with type-erasure in the sense that no type is needed at runtime. However, the unit element has to be present during runtime, because, although it does not contain information in itself, its sole existence is a proof that code depending on it behaves correctly. We developed this example of non-erasable coercions with incoherent polymorphism in Section 5.5.

A second definition of non-erasable coercions is to simulate a rich runtime over a poor runtime. Coercions would be erasable in the rich runtime while non-erasable in the poor one. An illustration of non-erasable coercions in this setting are record subtyping. A rich runtime may freely forget additional fields, while the poor runtime would need to copy the record value without its additional fields. For example, in the rich runtime `{name = "bob"; age = 23}` accepts type `{age : Int}`, while in the poor runtime it only has type `{name : String; age : Int}` and needs to be coerced with non-erasable code to have type `{age : Int}`. OCaml has actually both forms of coercions. The subtyping for variants and objects is using a rich runtime: coercions are erasable. And the subtyping for modules uses a poor runtime: coercions copy the module by η -expansion.

Finally, a loose definition is to allow any form of code as coercions. The typing implications of this definition is not obvious, but it may be a framework to define code inference by restraining the inferred code to be valid according to the coercion judgment. We might expect the main feature of such coercions to be η -expansion, because it follows the structure of types. The rule for coercion application would resemble the following typing rule:

$$\frac{\text{TERMCOER} \quad M \Rightarrow \Gamma, \Sigma \vdash a : \tau \quad G \Rightarrow \Gamma \vdash C : (\Sigma \vdash \tau) \triangleright \sigma}{G\langle M \rangle \Rightarrow \Gamma \vdash C[a] : \sigma}$$

This rule differs from its preceding version in Figure 4.10 by the presence of a computational context C in the coercion judgment. Similarly to the computational term a in the term judgment, we now say that the coercion G erases to the context C . And we can see that the erasure of the coercion application $G\langle M \rangle$ is now a context application $C[a]$. It is interesting to study restrictions of this general setting in the framework of implicit coercions (see Section 6.2.3): boring coercions could be inferred.

6.1.6 First-class coercions

First-class coercions is the possibility to pass coercions as arguments to functions and to return coercions as objects. A simple way to achieve this feature is to use incoherent existential types (see Section 6.1.2 and 5.5). This solution does not permit to reduce parts of terms relying on first-class coercions to preserve soundness.

We actually define first-class propositions, which imply first-class coercions, because coercions are propositions. We write $[P]$ the first-class type for the proposition P . It is a notation for $\Sigma(\alpha : 1 \mid P)1$, where 1 is the type containing only the unit term. When we receive a first-class proposition as argument we may unpack it to access the proposition in the constrained kind. Since we use the incoherent version of the unpack construct, its body cannot be reduced until we know that the proposition holds, which we know when the associated pack construct

is made accessible. We may also return such first-class proposition by packing. The incoherent packing construct has no reduction limitations.

One example of first-class propositions is to write a function of type $[\tau_1 \equiv \tau_2] \rightarrow \sigma$ which relies on the fact that τ_1 and τ_2 are equal to return a term of type σ .

6.1.7 Dependent types

A much more difficult extension would be to redefine the framework starting from the dependently typed λ -calculus instead of the STLC. In a dependently typed setting the return type of a term abstraction may mention its term argument. As a consequence, terms may appear in types. Or said otherwise, types may depend on terms, which is why dependent types are called so. The arrow type changes from $\tau \rightarrow \sigma$ to $\Pi(x : \tau)\sigma$ (which should not be confused with the incoherent polymorphic type $\Pi(\alpha : \kappa)\tau$), where the term argument x of type τ is bound in the return type σ . Rule `TERMLAM` shows this new typing for term abstraction. In rule `TERMAPP`, we can see that term application remembers its argument in its type. It is the return type σ of its function where the term argument variable x has been substituted with the actual argument b .

$$\frac{\text{TERMLAM} \quad a : \Gamma, (x : \tau) \vdash \sigma}{\lambda x a : \Gamma \vdash \Pi(x : \tau)\sigma} \qquad \frac{\text{TERMAPP} \quad a : \Gamma \vdash \Pi(x : \tau)\sigma \quad b : \Gamma \vdash \tau}{ab : \Gamma \vdash \sigma[x/b]}$$

The implications of such extension on the framework and the semantics are not negligible. The following works handle erasable contents in a dependent setting. However, these works do not take a coercion approach with a distinct composable judgment for erasable typing transformations. They reuse the term judgment by adding implicit rules as we usually see.

The Implicit Calculus of Constructions (ICC) defined in [24] can be seen as an extension of System F_η to dependent types. The presentation uses the η -expansion term judgment rule, where a has type $\tau \rightarrow \sigma$ if $\lambda x a x$ has type $\tau \rightarrow \sigma$ and x is not free in a . The subtyping judgment (between types) is derived from the term judgment. System F_{cc} extended with dependent types would extend ICC with coercion abstraction. Some major difficulties are expected when dealing with strong normalization, types (which are terms) inhabitation, and semantics. A work by Bruno Barras and Bruno Bernardo [6] gives an explicit version to a restriction of ICC. Large elimination is handled in [1].

Instead of extending the framework to handle dependent types, another solution is to simulate dependent type features by lifting terms to types and types to kinds, and using singleton types to link term values to their associated types. This solution is used in System `FC` (see Section 6.1.8).

6.1.8 Kind coercions

System `FC` has kind equality coercions [39]. The goal of such feature is to give GHC a type system with dependent type features. Actually, they achieve this by merging types and kinds into the same syntactical class. As a consequence, type equality coercions naturally extend to kind equality coercions. In System F_{cc} , studying kind coercions (a form of subkinding) seems natural, because kinds are sets of types and are thus naturally ordered. In particular the constrained kind $\{\alpha : \kappa \mid P\}$ is by definition a subkind of κ . But we also have that $\{\alpha : \kappa \mid P_1\}$ is a subkind of $\{\alpha : \kappa \mid P_2\}$ if P_1 implies P_2 .

Notice that we already have a notion of subkinding for coherence: a kind κ is a subkind of the kind κ' under the type environment Σ , if the coherence of κ under Σ implies the coherence of κ' under Σ . We can prove such propositions with rule `PROPSTR` where the conclusion proposition is $\Sigma; \emptyset; \emptyset \vdash \exists \kappa'$ and the kind coherence premise is $\Sigma \vdash \kappa$. This notion of subkinding for coherence is limited and cannot be used to rekind a type. However, we somehow use this idea of coherence propagation with the instance type σ' in the η -expansion rule of incoherent polymorphism `COERP1`.

There are two alternatives to study kind coercions. The first path is to keep types and kinds distinct and develop a new theory about kind coercions. Another path is to mimic `FC` and merge types and kinds and reuse the type coercions mechanism for kinds too. The first solution may look safer, and maybe easier to prove, but it may also imply more duplication of rules.

6.1.9 Intersection types

We could perhaps add binary intersection types as coercions following the approach of Wells on *branching types* [40]. The same way we extended subtyping from types to typings, we could extend branching types to *branching typings*, which would be trees of typings where leaves are usual typings and nodes are chunks of typing environments. We would have the following grammar for invariants, instead of the usual single typing $\Gamma \vdash \tau$.

$$\Phi ::= \Gamma \vdash \tau \mid \Gamma \langle \Phi \cap \Phi \rangle \quad \text{Branching typings}$$

For the intuition, we can define a function $[\cdot]$ from branching typings to sets of usual typings. We would have $[\Gamma \vdash \tau]$ be the singleton set containing the typing $\Gamma \vdash \tau$. And we would define $[\Gamma \langle \Phi_1 \cap \Phi_2 \rangle]$ as $\{\Gamma, \Gamma' \mid \Gamma' \in [\Phi_1] \cup [\Phi_2]\}$. Intuitively, if a term accepts the branching typing Φ , it means that it accepts all the typings in $[\Phi]$. In other words, a term is in the semantics $|\Phi|$ if it is in the intersection of the semantics of the elements of $[\Phi]$.

We would also need to modify the usual `STLC` typing rules to act simultaneously on all leaves. For instance, we would need to define $\Phi_1 \times \Phi_2$, where Φ_1 and Φ_2 are binary trees with the same structure where only the types at the leaves may change, as the same binary tree where leaves $\Gamma \vdash \tau_1$ and $\Gamma \vdash \tau_2$ become $\Gamma \vdash \tau_1 \times \tau_2$. Coercions are now between these branching typings. We also need to add coercions from $\tau \cap \sigma$ to τ and σ , and a coercion from the branching typing $\langle (\emptyset \vdash \tau) \cap (\emptyset \vdash \sigma) \rangle$ to the type $\tau \cap \sigma$.

6.1.10 Semantic consistency

In System `Fcc`, the coherent polymorphism abstraction rule is well-typed when the kind is inhabited. When the kind is a constrained kind $\{\alpha : \kappa \mid P\}$, the proposition P has to hold for at least one inhabitant of κ . Currently, this proof of coherence is done using syntactical and concrete witnesses. We could extend the logic to prove propositions with more semantic proofs, or we may delay proof of propositions to top-level and use some algorithm to check them. We can already model this last feature using polymorphic propositions and a top-level type environment Σ_0 with all the propositions we will ever use in the derivation. For example, when we need to prove $\Sigma; \Theta_0; \Theta_1 \vdash P$, we may add a top-level type binding of the form $(\alpha : 1 \mid \forall \Sigma P)$ in Σ_0 . Then, instead of proving the coherence of Σ_0 using the syntactical rules of Figure 5.7, we may prove its semantic interpretation using mathematics, or with an algorithm. This is actually the solution followed in `Constraint ML`.

6.1.11 Environment actions

The framework of coercions defined in Part II does not handle coercions that would act on the environment in other ways than bindings. For instance, the weakening coercion shrinks the current environment while usual constructs extend the environment.

If we were considering coercions in their generality and not only coercion extending environments, how should we present them? They should still be written with the judgment given in System F_t^p , namely $G \Rightarrow \Gamma \vdash (\Sigma \vdash \tau) \triangleright \sigma$, however, the meaning of the erasable environments Σ would change to an environment action. An environment action is the modification a coercion may do on its surrounding environment. For instance, in System F_t^p with erasable environments, the two environment actions we have are: type binding and coercion binding. However we all know that we may come up with more environment actions like: weakening, binding swapping, or environment shuffling.

As a consequence of this modification from erasable environment to environment actions, we need to modify the notion of simple concatenation we had Γ, Σ , to a more subtle environment actions composition $\Gamma + \Sigma$. For example, if we define $W\alpha$ the weakening binding for α , then the following environment actions application $\Gamma, \alpha, \Gamma' + W\alpha$ evaluates to Γ by removing the type binding α and all the following bindings that may depend on it. The term typing rule for coercions now becomes:

$$\frac{\text{TERM COER} \quad M \Rightarrow a : \Gamma + \Sigma \vdash \tau \quad G \Rightarrow \Gamma \vdash (\Sigma \vdash \tau) \triangleright \sigma}{G\langle M \rangle \Rightarrow a : \Gamma \vdash \sigma}$$

and the coercion typing rules for transitivity becomes:

$$\frac{\text{COERTRANS} \quad G_1 \Rightarrow \Gamma + \Sigma_2 \vdash (\Sigma_1 \vdash \tau_1) \triangleright \tau_2 \quad G_2 \Rightarrow \Gamma \vdash (\Sigma_2 \vdash \tau_2) \triangleright \tau_3}{G_1 \circ G_2 \Rightarrow \Gamma \vdash (\Sigma_2, \Sigma_1 \vdash \tau_1) \triangleright \tau_3}$$

where Σ_2, Σ_1 is the concatenation or sequencing of environment actions: we first do the actions of Σ_2 , followed by the actions of Σ_1 : $\Gamma + (\Sigma_2, \Sigma_1)$ evaluates like $(\Gamma + \Sigma_2) + \Sigma_1$.

Environment actions subsume the more intuitive notation $\Gamma \vdash (\Sigma_1 \vdash \tau_1) \triangleright (\Sigma_2 \vdash \tau_2)$. Indeed, this last coercion can be encoded as $\Gamma, \Sigma_2 \vdash (-\Sigma_2 + \Sigma_1 \vdash \tau_1) \triangleright \tau_2$, where $-\Sigma_2$ is the opposite action of Σ_2 . This coercion can be used in a similar way, because the typing $\Gamma, \Sigma_2, -\Sigma_2, \Sigma_1 \vdash \tau_1$ is actually equivalent to $\Gamma, \Sigma_1 \vdash \tau_1$.

6.1.12 Coercion reduction

When the coercion language gets bigger or when an intensive use of them is done, it becomes interesting to reduce them to make the typing derivations smaller. Coercion reduction may also be necessary to prove soundness (see Section 4.6.2) using the fact that coercions strongly normalize to a coercion value, which we can decompose.

In System F_t^p , our term reduction relation contains β and ι steps. We never reduce coercions directly, but only when applied to terms. The reason is that there is no need for doing so since our coercions can be erased. Moreover, it keeps the presentation simpler.

Still, introducing a reduction relation $G \rightsquigarrow_\gamma G$ between coercions themselves is possible, and could be interesting in other contexts. For example, coercions in System FC [38] are elaborated from the surface language, inferred from the constraint solver, and made bigger

by term optimizations. In all these cases, coercions may grow very large. Hence, reducing coercions is a way to make them smaller.

We briefly describe how to add coercion reduction in our language. Since we use strong reduction, coercion evaluation contexts are all coercion contexts:

$$\begin{array}{c} \text{REDIGAMMA} \\ \frac{G_1 \rightsquigarrow_\gamma G_2}{G_1 \langle M \rangle \rightsquigarrow_\iota G_2 \langle M \rangle} \end{array} \qquad \begin{array}{c} \text{REDGCTX} \\ \frac{G_1 \rightsquigarrow_\gamma G_2}{F[G_1] \rightsquigarrow_\gamma F[G_2]} \end{array}$$

Rule REDIGAMMA is the ι -reduction linking γ -reduction to ι -reduction. If a coercion reduces, then its application to a term reduces too. Rule REDGCTX is the context rule for γ -reduction. Then, some obvious reduction rules are pushing transitivity down, pulling reflexivity up, and reducing coercion redexes:

$$\begin{array}{c} \text{REDGETAARR} \\ (G_1 \xrightarrow{\tau} G_2) \circ (G'_1 \xrightarrow{\tau'} G'_2) \rightsquigarrow_\gamma (G'_1 \circ G_1) \xrightarrow{\tau} (G_2 \circ G'_2) \end{array}$$

$$\begin{array}{c} \text{REDGETAPROD} \\ (G_1 \times G_2) \circ (G'_1 \times G'_2) \rightsquigarrow_\gamma (G_1 \circ G'_1) \times (G_2 \circ G'_2) \end{array} \qquad \begin{array}{c} \text{REDGFOLD} \\ \text{unfold} \circ \text{fold}_{\mu\alpha.\tau} \rightsquigarrow_\gamma \diamond \end{array}$$

Rules REDGETAARR, REDGETAPROD, and REDGFOLD are simple transposition of their ι counterparts (which could then be removed). These reduction rules are sound. They are also incomplete: other rules, which are not derivable could also be added—but they require other operations on coercions that have not been defined yet. For instance, all ι -reduction rules involving no untyped constructs on their left-hand side could be transformed into the combination of an equivalent γ -reduction rule and rule REDIGAMMA. For instance, one remaining case is rule REDTYPE. If we had a coercion for explicit type substitution, we would be able to have the following γ -reduction rule:

$$\begin{array}{c} \text{REDGTYPE} \\ \bullet \tau \circ \Lambda \alpha \rightsquigarrow_\gamma [\alpha/\tau] \end{array}$$

Notice that one must now prove the termination of γ -reduction rules since they are included in ι -reduction through rule REDIGAMMA.

6.1.13 Side effects

Extending the λ -calculus with side effects is a natural question since real-world programs do have side effects. The main modification is to add a notion of memory and how programs interact with it. Since this memory has a semantic meaning, invariants have to be stated according to the memory. Invariants are no longer a pair of an environment and a type, but a triplet of an environment, a memory invariant, and a type. It should be possible to adapt the step-indexed semantics according to existing work for weak reduction [4]. We don't expect existing problems (like value restriction) to disappear with strong reduction.

6.1.14 Dead code

Although incoherent polymorphism is required to express GADTs (see Section 5.5), we might still want to discriminate between potentially inconsistent branches and always inconsistent branches. A branch is potentially inconsistent if its coherence depends on the instantiation of

its environment, while a branch is always inconsistent if it is incoherent for all its environment instantiations. For example, the kind $\{\beta : 1 \mid \alpha \equiv \mathbf{Int}\}$ in the environment \emptyset, α , is potentially incoherent. It is coherent if the type variable α is instantiated with \mathbf{Int} , but it is incoherent if the type variable α is instantiated with \mathbf{Bool} . Let's now consider the kind $\{\beta : 1 \mid \mathbf{Bool} \equiv \mathbf{Int}\}$ in the empty environment. It is always incoherent, and we might want to either reject terms generalizing on this kind or warn the user that he is writing dead code. Such type abstraction is actually not instantiable.

In pedagogical systems [10], an abstraction (such as type abstraction) is *useful*, if its domain (such as kinds in the case of type abstraction) is inhabited for some instantiation of the environment. This condition implies that the abstraction is not dead code and can be instantiated. We can use this idea to define an alternative version of incoherent polymorphism, by strengthening its condition from the kind to be well-formed to the kind to be coherent for some instantiation of the environment. This useful incoherent polymorphic type $\overline{\Pi}(\alpha : \kappa)\tau$ would have the following introduction rule:

$$\frac{\text{TERMGENUSEFUL} \quad \exists \Sigma \kappa \quad a : \Sigma, (\alpha : \kappa); \Gamma \vdash \tau \quad \alpha \notin \text{fv}(\Gamma)}{\partial a : \Sigma; \Gamma \vdash \overline{\Pi}(\alpha : \kappa)\tau}$$

The premise $\exists \Sigma, \kappa$ asks the kind κ to be coherent for some instantiation of Σ . This rule is in the term judgment, because the usefulness does not imply consistency and the abstraction cannot be erasable. We can reuse the existing type abstraction, which we already use for incoherent polymorphism, because we can always tell them apart by typing. The other rules for useful incoherent polymorphic types are identical to the incoherent polymorphic versions.

Such extension would tell us that we cannot reduce under the useful incoherent abstraction now, but that it is possible to reach this part of the code for some instantiation of the environment.

Unfortunately, this property is not stable by reduction, because reduction could instantiate the environment making the kind incoherent for this particular instantiation. If one want this dead code property to be stable by reduction, he would need to have some form of coherent polymorphism, which we already have.

6.2 Related work

To the best of our knowledge there is no previous work considering coercions as an inclusion of typings. However, the use of coercions to study features of type system is not at all new. This section presents how coercions have been used to study subtyping, GADTs, recursive types, etc.

Our work on step-indices is greatly inspired from existing works in weak reduction settings. We actually merged the typical semantic proofs of System F in a strong reduction setting, with the idea of fuel to limit the reduction of terms.

6.2.1 System $F_{<}$:

Record subtyping in System $F_{<}$ may be compiled away into records without subtyping in plain System F by inserting coercions with computational content [9] that change the representation of records whenever subtyping is used. Since these coercions are not erasable and can be inserted in different ways, the soundness of the approach depends on a coherence result to

show that the semantics of the translation does not actually depend on the places where coercions are inserted.

Another method for eliminating subtyping has been used by Crary [12]: bounded polymorphism $\forall(\alpha \triangleright \tau) \sigma$ is compiled away into an intersection type $\forall \alpha \sigma[\alpha/\alpha \cap \tau]$ while intersection types are themselves encoded with explicit erasable coercions. This directly relates to our work by their canonization, which is similar to our ι -reduction, and their use of bisimulation up to canonization to show erasability of coercions. Of course, the languages are different, as we do not consider intersection types while they do have neither coercion abstraction nor distributivity and only consider call-by-value reduction.

6.2.2 System FC

In System FC [35] (the core language of the Glasgow Haskell Compiler), coercions are bidirectional: they are proofs of equality instead of proofs of inclusion. As a consequence they are structural: when two types are equal according to an equality coercion, then their head type constructors are the same. This permits drastic simplifications and in particular to handle push much more easily. Notice that consistency already asks the computational head type constructors of coercions to be equal. In FC, this condition applies to all type constructors, even the erasable ones. As a consequence, the coercion $(\forall \alpha \tau \rightarrow \sigma) = (\tau \rightarrow \forall \alpha \sigma)$, when α is not free in τ , is ill-formed, even though these two types are semantically equal: $(\forall \alpha \tau \rightarrow \sigma) \equiv (\tau \rightarrow \forall \alpha \sigma)$ holds in System F_t^p and System F_{cc} .

Coercion abstractions are not fully erasable in System FC, but they are zero-bit, because they only have to do with typings and do not modify the computational content of the term they are applied to. They are not erased because they have to block the reduction, as coercion abstraction in FC is incoherent. See Section 5.5 to see why incoherent coercion are useful and how we extend System F_{cc} with incoherent polymorphism. Only top-level coercion axioms are checked for consistency, because we have to reduce under them.

In System FC, coercions are not as powerful as in other type systems, since they are only equality and structural, however they permit to define a number of derivable type system features such as GADTs, type families, type synonyms, etc.

System FC actually has an heterogeneous equality for coercions and provides an additional notion of kind equality. This feature permits to lift ([41] and [39]) all sorts of data types to the type level and use them to index singleton types in order to simulate dependent type mechanisms.

6.2.3 Implicit Coercions

The notion of implicit objects (implicit terms, implicit coercions, etc.) is a feature of surface languages and in particular inference. It is orthogonal to the notion of erasability. On the one hand, an object is *erasable* if it is present during typing but absent at runtime. For instance, types and coercions are erasable in System F_{cc} . On the other hand, an object is *implicit* if it is absent in the source code but may be present at runtime. Most of the time, inference elaborates implicit objects from the surface language to the kernel language. So the study of implicit objects is the study of inference.

In Coq, coercions are not erasable but they are implicit. They are not added to extend the expressiveness of the language, but to lighten the source language. They permit to implicitly

coerce a term of a particular type to another type, by giving particular computational functions which may be used in this manner.

Implicit coercions may also be used to express a wide variety of practical features, from dynamic software updating to provenance tracking [37]. However, the more powerful the inference of coercions, the greater the possibility of several semantically different elaborations. As a consequence, trade-offs between expressiveness and ambiguity have been studied. This notion of ambiguity can be solved by proving that all rewriting of the source term into a kernel term with coercions have the same semantics, no matter how they are introduced. This path is followed by [36] to ease programming with monads. The binds, units, and monad-to-monad morphisms necessary when programming in a monadic-style, are inferred based on types. This idea is called coherence in [7], which uses implicit coercions to control logical properties in pure type systems.

In coercive subtyping [21], a type τ is a subtype of σ whenever there is a unique coercion from τ to σ . Coercive subtyping extends an existing type theory T with a subtyping judgment C to form a new type system $T[C]$. The coherence of coercions comes from the equivalence of $T[C]$ with $T[C]^*$, which is a version of $T[C]$ where the position of coercions are marked. As a consequence, coercive subtyping lightens the source code without ambiguity. This work is done in a dependent type setting and solves the difficulty of Σ -types [20], which is that the congruence on the first component of a Σ -type and the first projection of a Σ -type form a set of incoherent coercions.

This approach of coercions as implicit objects is orthogonal to our work and may thus be used at the same time, to get the most of both worlds. Implicit coercion inference happens from the source language to the kernel language, while erasable coercions are a property of the erasure of the kernel language, which is independent.

6.2.4 Step-indices

Numerous works have been done on step-indices ([4], [15]), and all that we know of are in a weak-reduction setting. The reason is probably that step indices are usually used to give a semantics to programming languages with side effects, which usually come with a weak reduction strategy. In all these works there is a distinction between the values of type τ , written $\mathcal{V}[\tau]$, and expressions of type τ , written $\mathcal{E}[\tau]$ and defined by the terms having their normal form in $\mathcal{V}[\tau]$ with the remaining fuel. The set of values of arrow type $\tau \rightarrow \sigma$ is then the set of term abstractions $\lambda x a$ with index k such that for all indices j smaller than k and for all values v at j in $\mathcal{V}[\tau]$, the substitution $a[x/v]$ is $\mathcal{E}[\sigma]$ at j . This definition is not stable by reduction if we reduce under the abstraction. However, step-indices are not only used for side-effects, they are also used for recursive types, which makes perfect sense in a language with strong reduction, as we do. This is why we use them.

6.3 Applications

There are several applications to this work, although most of them would be more practical with some of the extensions discussed in Section 6.1. However, we can see System F_{cc} as a good starting point for these applications.

Designing features that are easy to merge Our initial theoretical motivation was to express type systems features in a unique framework and in an orthogonal way. When type

system features are written as coercions, they can easily be composed with pre-existing type system features. Besides the ease of merging, studying type system features as coercions naturally gives the most out of the fusion of two features. The typical example is for System $F_{<}$: which merges η -expansion (called subtyping) with upper bounded polymorphism. Merging these two features using coercions naturally gives the distributivity rule which is not present in System $F_{<}$.

Extending FC with subtyping A more practical motivation to study coercions in a general setting, was the use of type equality coercions in FC to express useful surface features such as GADTs, type families, and others. There are several main differences between System F_{cc} and FC. On the one hand, our type system only has an implicit version, although we discuss a possible explicit version in Section 5.7. On the other hand, our type system features subtyping. A very important and potentially difficult extension to consider are side-effects (Section 6.1.13). A less important but potentially difficult extension are kind coercions (Section 6.1.8).

The other differences between both systems are more anecdotal. Zero-bit coercion abstractions of FC may be added using incoherent polymorphism (Section 5.5). Type-level functions of FC may be added as explained in Section 6.1.3. Type equality coercions are simply erasable isomorphisms: τ is equal to σ if there is a coercion from τ to σ and reciprocally. Higher-kind recursive types are required to define user data-types (Section 6.1.4). Finally, coercion reduction (see Section 6.1.12) may be considered to optimize the compiler, because surface language elaboration, constraint solving, and code optimization may build large coercions and slow down the compilation phase.

Merging inference systems If two type systems with inference rely on the same programming language, and both type systems can be projected into a common kernel type system, then a program may use one or the other inference system for each of its compilation unit, as long as the interfaces of the compilation units are in the intersection of both type systems. A concrete example would be to use MLF and Constraint ML with ML types for interfaces.

MLF and Constraint ML are surface type systems with powerful inference mechanism. Merging the two inference systems may be a very hard task. However, if one only needs to write some parts of his program using the MLF inference system and other parts with the Constraint ML system, he can partition his compilation units according to the inference system he wishes to use for them. Each compilation unit should have an interface in ML, which is a common sublanguage of MLF and Constraint ML.

Suppose we have a module M written in MLF with an interface in ML and a module N written in Constraint ML with an interface in ML. As we showed in Section 5.4, MLF and Constraint ML are subsumed in System F_{cc} . So, once M and N have been inferred in their respective language with the other module interface, we can translate their typing derivations to System F_{cc} and then link them to get the final program. The linking works correctly since the interface types are the same.

Chapter 7

Conclusion

We presented two type systems, System F_t^p and System F_{cc} , based on a coercion framework. Coercions are an extension of subtyping from type ordering to typing ordering. The particularity of the coercion framework is that all type system features are expressed as coercions. The underlying programming language we use is the λ -calculus extended with pairs. The main object of study is coercion abstraction.

Both type systems are sound. They are also strongly normalizing after the removal of recursive types. A type system is sound (resp. strongly normalizing) when all well-typed terms are sound (resp. strongly normalizing). A term is sound when it cannot reach an error state. A term is strongly normalizing when it cannot indefinitely reduce. When these properties hold in a strong reduction setting, they also hold in all usual reduction strategies (weak reduction, call by value, call by need, etc.). We thus described our type systems and proved these properties in a strong reduction setting.

Coercions are erasable typing transformations in the implicit version and erasable contexts in the explicit version. System F_t^p has an implicit and explicit version, while System F_{cc} only has an implicit version. With an implicit version, coercions are naturally erasable since they are not present in the term. With an explicit version, coercions are erasable if they satisfy the bisimulation lemma, which tells that coercions should not introduce, remove, or block computational steps. System F_t^p satisfies the bisimulation lemma.

Both type systems feature polymorphism, η -expansion, and coercion abstraction as coercions. The main difference is that coercion abstraction is restricted to parametric coercions in System F_t^p , while it is unrestricted in System F_{cc} . Parametric coercions are coercions where either the argument type or the return type is an abstract type. Parametric coercion abstraction can actually be seen as an extension from polymorphism to bounded polymorphism. In System F_{cc} , we also extend polymorphism by extending the kinds on which abstract types may range, with a particular kind constraining the abstract type to satisfy a proposition, which may be a conjunction of coercions for instance. The η -expansion feature gives both type systems congruence rules for computational types. This permits to use coercions deeply in a type according to the variance of the path. This feature is at the root of subtyping and comes from System F_η .

Both type systems subsume System F_η , MLF, and System $F_{<}$. System F_{cc} additionally subsumes Constraint ML. As a consequence, MLF and Constraint ML which have particularly good inference mechanisms, can be seen as surface languages for System F_{cc} . The main missing feature, to describe OCaml core type system and FC (GHC core type system), is side effects.

The remaining features are more anecdotal to study.

Another main difference with **FC**, is that **FC** has an explicit version with the subject reduction lemma, while only System F_t^p does. System F_{cc} is more expressive than System F_t^p , since it features unrestricted coercion abstraction and unrestricted recursive types. But it only has an implicit version without subject reduction. Restoring subject reduction and defining an explicit version would require to study coercion decomposition and the push reduction rule. The main difficulty is the proof of consistency which is not completely clear to hold in our semantics. Moving polymorphism from the coercion language to the term language may be a solution to restore subject reduction and the explicit version at the cost of losing deep type and coercion abstraction and instantiation.

A last difference with **FC**, is that **FC** has top-level coherent coercion abstraction and local incoherent coercion abstraction. In System F_{cc} , we permit also local coherent coercion abstraction by asking the kind of the type abstraction to be inhabited. We also permit incoherent coercion abstraction as an extension, which requires blocking abstraction at the computational level: such as unit abstraction or weak lambdas.

A final contribution of this thesis is an adaptation of step-indexed techniques to prove soundness in a strong reduction setting. Step-indexed techniques are usually studied in a weak reduction setting due to the presence of side effects in the considered languages. However, the definition they use for the arrow type does not work for strong reduction semantics. The reason comes from the fact that indices, representing the fuel of the terms, are external to the terms. As a consequence, substitution and reduction do not permute. We restore this commutation by internalizing the notion of indices inside terms.

Bibliography

- [1] Andreas Abel and Gabriel Scherer. On irrelevance and algorithmic equality in predicative type theory. *Logical Methods in Computer Science*, 8(1):1–36, 2012. TYPES’10 special issue.
- [2] Beniamino Accattoli and Delia Kesner. The structural *lambda*-calculus. In Anuj Dawar and Helmut Veith, editors, *CSL*, volume 6247 of *Lecture Notes in Computer Science*, pages 381–395. Springer, 2010.
- [3] Roberto Amadio and Luca Cardelli. Subtyping recursive types. *ACM TRANSACTIONS ON PROGRAMMING LANGUAGES AND SYSTEMS*, 15(4):575–631, 1993.
- [4] Andrew W. Appel and David McAllester. An indexed model of recursive types for foundational proof-carrying code. *ACM Transactions on Programming Languages and Systems.*, 23(5), September 2001.
- [5] Paolo Baldan, Giorgio Ghelli, and Alessandra Raffaetà. Basic theory of F-bounded quantification. *Inf. Comput.*, 153:173–237, September 1999.
- [6] Bruno Barras and Bruno Bernardo. The implicit calculus of constructions as a programming language with dependent types. In *Proceedings of the Theory and practice of software, 11th international conference on Foundations of software science and computational structures*, FOSSACS’08/ETAPS’08, pages 365–379, Berlin, Heidelberg, 2008. Springer-Verlag.
- [7] Gilles Barthe. Implicit coercions in type systems. In Stefano Berardi and Mario Coppo, editors, *Types for Proofs and Programs*, volume 1158 of *Lecture Notes in Computer Science*, pages 1–15. Springer Berlin Heidelberg, 1996.
- [8] Tomasz Blanc, Jean-Jacques Lévy, and Luc Maranget. Sharing in the weak lambda-calculus. In Aart Middeldorp, Vincent Oostrom, Femke Raamsdonk, and Roel Vrijer, editors, *Processes, Terms and Cycles: Steps on the Road to Infinity*, volume 3838 of *Lecture Notes in Computer Science*, pages 70–87. Springer Berlin Heidelberg, 2005.
- [9] Val Breazu-Tannen, Thierry Coquand, Carl Gunter, and Andre Scedrov. Inheritance as implicit coercion. *Information and Computation*, 93:172–221, 1991.
- [10] Loïc Colson and Vincent Demange. Investigations on a pedagogical calculus of constructions. *CoRR*, abs/1203.3568, 2012.
- [11] The Coq Proof Assistant <http://coq.inria.fr/what-is-coq>.

- [12] Karl Crary. Typed compilation of inclusive subtyping. In *Proceedings of the fifth ACM SIGPLAN international conference on Functional programming (ICFP)*, pages 68–81, New York, NY, USA, 2000. ACM.
- [13] Julien Cretin and Didier Rémy. On the power of coercion abstraction. In *Proceedings of the 39th annual ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, POPL '12, pages 361–372, New York, NY, USA, 2012. ACM.
- [14] Kevin Donnelly and Hongwei Xi. A formalization of strong normalization for simply-typed lambda-calculus and system F.
- [15] Derek Dreyer, Amal Ahmed, and Lars Birkedal. Logical step-indexed logical relations. In *Proceedings of the 2009 24th Annual IEEE Symposium on Logic In Computer Science*, LICS '09, pages 71–80, Washington, DC, USA, 2009. IEEE Computer Society.
- [16] Jean H. Gallier and Jean H. Gallier. On Girard’s candidats de réductibilité. In *Logic and Computer Science*. North, 1990.
- [17] The Haskell Programming Language <http://www.haskell.org/haskellwiki/Haskell>.
- [18] Didier Le Botlan and Didier Rémy. Mlf: raising ml to the power of system f. In *Proceedings of the eighth ACM SIGPLAN international conference on Functional programming*, ICFP '03, pages 27–38, New York, NY, USA, 2003. ACM.
- [19] Didier Le Botlan and Didier Rémy. Recasting MLF. *Information and Computation*, 207(6):726–785, 2009.
- [20] Yong Luo and Zhaohui Luo. Combining incoherent coercions for sigma-types. In Stefano Berardi, Mario Coppo, and Ferruccio Damiani, editors, *TYPES*, volume 3085 of *Lecture Notes in Computer Science*, pages 276–292. Springer, 2003.
- [21] Z. Luo, S. Soloviev, and T. Xue. Coercive subtyping: Theory and implementation. *Inf. Comput.*, 223:18–42, February 2013.
- [22] Jean-Jacques Lévy and Luc Maranget. Explicit substitutions and programming languages. In C.Pandu Rangan, V. Raman, and R. Ramanujam, editors, *Foundations of Software Technology and Theoretical Computer Science*, volume 1738 of *Lecture Notes in Computer Science*, pages 181–200. Springer Berlin Heidelberg, 1999.
- [23] Ralph Matthes. Monotone fixed-point types and strong normalization. In *Proceedings of the 12th International Workshop on Computer Science Logic*, pages 298–312, London, UK, UK, 1999. Springer-Verlag.
- [24] Alexandre Miquel. The implicit calculus of constructions: extending pure type systems with an intersection type binder and subtyping. In *Proceedings of the 5th international conference on Typed lambda calculi and applications*, TLCA'01, pages 344–359, Berlin, Heidelberg, 2001. Springer-Verlag.
- [25] John C. Mitchell. Polymorphic type inference and containment. *Information and Computation*, 2/3(76):211–249, 1988.
- [26] The OCaml Programming Language <http://caml.inria.fr/ocaml/index.en.html>.

- [27] Martin Odersky, Martin Sulzmann, and Martin Wehr. Type inference with constrained types. *Theor. Pract. Object Syst.*, 5(1):35–55, January 1999.
- [28] Jens Palsberg and Patrick O’Keefe. A type system equivalent to flow analysis. *ACM Trans. Program. Lang. Syst.*, 17(4):576–599, July 1995.
- [29] Benjamin C. Pierce. *Types and programming languages*. MIT Press, Cambridge, MA, USA, 2002.
- [30] Benjamin C. Pierce. *Advanced Topics in Types and Programming Languages*. The MIT Press, 2004.
- [31] Francois Pottier. Simplifying subtyping constraints. In *In Proceedings of the 1996 ACM SIGPLAN International Conference on Functional Programming*, pages 122–133. ACM Press, 1996.
- [32] Didier Rémy and Boris Yakobowski. A church-style intermediate language for mlf. In Matthias Blume, Naoki Kobayashi, and German Vidal, editors, *Functional and Logic Programming*, volume 6009 of *Lecture Notes in Computer Science*, pages 24–39. Springer Berlin / Heidelberg, 2010.
- [33] Colin Riba. On the stability by union of reducibility candidates. In *Proceedings of the 10th international conference on Foundations of software science and computational structures*, FOSSACS’07, pages 317–331, Berlin, Heidelberg, 2007. Springer-Verlag.
- [34] Martin Steffen. Polarized higher-order subtyping, 1997.
- [35] Martin Sulzmann, Manuel M. T. Chakravarty, Simon Peyton Jones, and Kevin Donnelly. System F with type equality coercions. In *Proceedings of the 2007 ACM SIGPLAN international workshop on Types in languages design and implementation*, TLDI ’07, pages 53–66, New York, NY, USA, 2007. ACM.
- [36] Nikhil Swamy, Nataliya Guts, Daan Leijen, and Michael Hicks. Lightweight monadic programming in ML. In *Proceedings of the 16th ACM SIGPLAN international conference on Functional programming*, ICFP ’11, pages 15–27, New York, NY, USA, 2011. ACM.
- [37] Nikhil Swamy, Michael Hicks, and Gavin M. Bierman. A theory of typed coercions and its applications. In *Proceedings of the 14th ACM SIGPLAN international conference on Functional programming*, ICFP ’09, pages 329–340, New York, NY, USA, 2009. ACM.
- [38] Dimitrios Vytiniotis and Simon L. Peyton Jones. Evidence normalization in system fc (invited talk). In Femke van Raamsdonk, editor, *RTA*, volume 21 of *LIPICs*, pages 20–38. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2013.
- [39] Stephanie Weirich, Justin Hsu, and Richard A. Eisenberg. System fc with explicit kind equality. In *Proceedings of the 18th ACM SIGPLAN international conference on Functional programming*, ICFP ’13, pages 275–286, New York, NY, USA, 2013. ACM.
- [40] Joe B. Wells and Christian Haack. Branching types. In *Proc. of the European Symposium On Programming Languages and Systems*, 2002.

- [41] Brent A. Yorgey, Stephanie Weirich, Julien Cretin, Simon Peyton Jones, Dimitrios Vytiniotis, and José Pedro Magalhães. Giving Haskell a promotion. In *Proceedings of the 8th ACM SIGPLAN workshop on Types in language design and implementation*, TLDI '12, pages 53–66, New York, NY, USA, 2012. ACM.

List of Figures

2.1	Syntax of the λ -calculus	23
2.2	Reduction relation	24
2.3	Notations	24
3.1	STLC syntax	34
3.2	STLC notations	35
3.3	STLC reduction relation	35
3.4	STLC term judgment relation	36
3.5	STLC well formedness relations	36
3.6	STLC drop function	37
3.7	System F syntax	40
3.8	System F notations	40
3.9	System F reduction rules	41
3.10	System F term judgment relation	41
3.11	System F well-formedness relations	42
3.12	System F drop function	42
3.13	System F_{rec} syntax	44
3.14	System F_{rec} notations	45
3.15	System F_{rec} reduction rules	45
3.16	System F_{rec} term judgment relation	46
3.17	System F_{rec} well-foundedness judgment relation	46
3.18	System F_{rec} well-formedness relations	47
3.19	System F_{η} syntax	48
3.20	System F_{η} notations	49
3.21	System F_{η} reduction rules	49
3.22	System F_{η} term judgment relation	50
3.23	System F_{η} containment judgment relation	51
3.24	System F_{η} well-formedness relations	52
3.25	MLF syntax	54
3.26	MLF notations	54
3.27	MLF reduction rules	55
3.28	MLF update function	55
3.29	MLF term judgment relation	56
3.30	MLF instance judgment relation	56
3.31	MLF well-formedness relations	57
3.32	System $F_{<}$ syntax	58
3.33	System $F_{<}$ notations	59

3.34	System $F_{<}$: reduction rules	59
3.35	System $F_{<}$: term judgment relation	60
3.36	System $F_{<}$: subtyping judgment relation	61
3.37	System $F_{<}$: well-formedness relations	62
3.38	Constraint ML syntax	63
3.39	Constraint ML term judgment relation	63
4.1	Base system syntax	70
4.2	Base system notations	70
4.3	Base system reduction rules	72
4.4	Base system term judgment relation	72
4.5	Base system coercion judgment relation	73
4.6	Base system well-formedness relations	74
4.7	System F_t^p syntax	81
4.8	System F_t^p notations	82
4.9	System F_t^p reduction rules	82
4.10	System F_t^p term judgment relation	85
4.11	System F_t^p coercion judgment relation	86
4.12	System F_t^p well-formedness relations	88
4.13	System F_t^p drop function	89
4.14	System F_t^p term reification function	90
4.15	System F_t^p type reification function	90
4.16	System F_t^p coercion reification function	91
4.17	System F_t^p environment reification function	91
4.18	Push	107
5.1	System F_{cc} syntax	112
5.2	System F_{cc} well-foundedness judgment relation	113
5.3	System F_{cc} kind judgment relation	114
5.4	System F_{cc} type judgment relation	114
5.5	System F_{cc} proposition judgment relation	115
5.6	System F_{cc} coercion judgment relation	117
5.7	System F_{cc} type environment judgment relation	118
5.8	System F_{cc} well-formedness relation	118
5.9	System F_{cc} term environment judgment relation	119
5.10	System F_{cc} term judgment relation	119
5.11	Indexed Calculus syntax	121
5.12	Indexed Calculus notations	121
5.13	Lower function	122
5.14	Indexed Calculus reduction relation	123
5.15	Lifting integer predicates to indexed terms	124
5.16	Lifting of a binary predicate \star on indices to terms	124
5.17	Kind interpretation	131
5.18	Type interpretation	131
5.19	Proposition interpretation	132
5.20	Environments interpretation	133
5.21	Fill function	134

5.22	System F_t^p type translation function	138
5.23	System F_t^p environment translation function	138

Index

Notations, *see* notations

$|\cdot|$, *see* interpretation

$\lfloor e \rfloor$, *see* drop

$\lfloor R \rfloor$, *see* drop

$\lfloor e \rfloor_k$, *see* lower

$\lceil a \rceil^k$, *see* fill

$R \downarrow$, *see* interior

$(R \rightsquigarrow)$, *see* contraction

$(\rightsquigarrow^* R)$, *see* expansion

$\diamond R$, *see* expansion-closure

$\langle R \rangle_k$, *see* approximation

$R \rightarrow S$, *see* arrow operator

$R \times S$, *see* product operator

$\forall F$, *see* intersection operator

μF , *see* recursive operator

$G \models S$, *see* semantic judgment

Δ , *see* head normal form

∇ , *see* neutral

Ω , *see* error

\mathcal{U} , *see* valid

r , *see* error

\mathcal{S} , *see* sound

$e \leq k$, *see* lift

$e > 0$, *see* lift

approximation, 127

arrow operator, 125, 149

bisimulation, 37, 42, 93

bottom type, 53, 77

coercion, 64, 72, 112

abstraction, 104, 111

decomposition, 106, 148

first-class, 155

implicit, 161

non-erasable, 154

reduction, 106, 158

coherence, 111, 114, 117, 143, 160

coinduction, 140

computational type, 74

confluence, 28, 38, 43, 92

congruence, 75

consistency, 93, 106, 111, 157

Constraint ML, **63**, 139

contraction, 125

curryfication, 29

dependent type, 156

distributivity, 51, 77, 103

drop, 37, 41, 123

equality, 112

equi-recursive type, 47, 140

equivalence, 37, 42, 53, 89

erasable type, 74

error, 24, 25, 122

η -expansion, 48, 76, 86, 116, 143

existential type, 144, 151

expansion, 120

expansion-closure, 125

explicit type system, 33

expressivity, 99, 135

extraction, 38, 43, 53, 96, 130

fill, 134

fuel, *see* step-index

GADT, 141

head normal form, 122

head normal form, 120

ICC, *see* Implicit Calculus of Constructions

Implicit Calculus of Constructions, 156

implicit type system, 33

Indexed Calculus, 121

inference, 53, 63

interior, 125

interpretation, 130

intersection operator, 127

iso-recursive type, 47, 140
 isomorphism, 137

 kind, 111

 Lambda Calculus, *see* λ -calculus
 λ -calculus, 21
 lift, 124
 lower, 122

MLF, **53**, 78, 102

 neutral, 120, 122
 non-expansive, *see* well-foundedness
 notations, 135

 omega term, 25, 47

 polymorphism, 39, 75
 bounded, 53, 58, 78, 79
 coherent, 111, 117, 135
 incoherent, 141
 pretype, 125
 product operator, 125
 progress, 39, 44, 98, 121
 proposition, 112
 push, 106, 149

 recursive bound, 56, 78–80, 88, 102
 recursive operator, 127
 recursive type, 44, 111, 127, 140, 153
 reification, 90

 semantic judgment, 128
 semantics, 120
 Simply Typed Lambda Calculus, *see* STLC
 sound, 120, 123
 soundness, 29, 39, 44, 95, 99, 130, 134
 step-index, 111, 121, 127, 162
STLC, **34**, 53, 71
 semantics, 128
 strong normalization, *see* termination
 subject reduction, 38, 44, 98, 148
 subtyping, 48, 58, 160
System F, **39**, 44, 48, 53, 58, 90, 99
System F_{rec}, **44**, 140
System F _{η} , **48**, 100
System F_<, **58**, 79, 103
System F_t^p, **69**, 137

System F_{cc}, **111**
 termination, 29, 38, 43, 90, 98, 123, 135
 top type, 58, 78
 type, 125
 type equality, *see* equality
 type system, 33

 uniqueness, 37, 41, 52, 89

 valid, 25, 122
 variance, 153

 weak reduction, 145
 weakening, 73, 85, 116
 well-founded, *see* well-foundedness
 well-foundedness, 45, 112
 semantics, 127