

Full reduction in the face of absurdity

Gabriel Scherer, Didier Rémy

► **To cite this version:**

Gabriel Scherer, Didier Rémy. Full reduction in the face of absurdity. [Research Report] INRIA. 2014.
<hal-01093910>

HAL Id: hal-01093910

<https://hal.inria.fr/hal-01093910>

Submitted on 11 Dec 2014

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

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

Full reduction in the face of absurdity

Gabriel Scherer¹ and Didier Rémy¹

INRIA,
{gabriel.scherer,didier.remy}@inria.fr

Abstract. Core calculi that model the essence of computations use full reduction semantics to be built on solid grounds. Expressive type systems for these calculi may use propositions to refine the notion of types, which allows abstraction over possibly inconsistent hypotheses. To preserve type soundness, reduction must then be delayed until logical hypotheses on which the computation depends have been proved consistent. When logical information is explicit inside terms, proposition variables delay the evaluation by construction. However, logical hypotheses may be left implicit, for the user’s convenience in a surface language or because they have been erased prior to computation in an internal language. It then becomes difficult to track the dependencies of computations over possibly inconsistent hypotheses.

We propose an expressive type system with implicit coercions, consistent and inconsistent abstraction over coercions, and *assumption hiding*, which provides a fine-grained control of dependencies between computations and the logical hypotheses they depend on. Assumption hiding opens a continuum between explicit and implicit use of hypotheses, and restores confluence when full and weak reductions are mixed.

Extended version For reasons of page limits, the proofs have been omitted from this version submitted for review. A full version of the present article, with additional remarks and all proofs, is also available electronically.¹

1 Introduction

The Curry-Howard isomorphism trained generations of statically-typed-language designers to be able to instantly switch their point of view from *programs* to *proof terms*, and from *types* to *logic statements*. Proof assistants based on type theory let us use our functional programming intuitions to *program* proofs. One example of the merits of such a re-unification is the strikingly simple and natural treatment of *axioms* in the functional languages of those assistants: assuming an axiom P is just abstracting over a variable $(x : P)$ of the corresponding type, and using this assumption is done by applying or pattern matching this bound variable x . These languages generally allow *full reduction*, in particular reducing under unapplied λ -abstractions. For example, a Coq program abstracting over

¹ At <http://gallium.inria.fr/~remy/coercions/>

an axiom P is of the form $\lambda(x : P) a$, where a may be computed as usual, but the reduction of subterms depending on x will be blocked.

There is a subtle but important contrast with how logical assumptions have been dealt with in languages designed mostly for programming rather than proving, such as `ML` or `Haskell`. Typical examples are the reasoning on type equalities in the `ML` module system, or in Generalized Algebraic Data Types (GADTs). Consider the following example, that implements application up to type equality, given in OCaml-like syntax:

```
type (., _) eq = Refl : (α, α) eq
let apply : ∀α1α2β. (α1 → β) → α2 → (α1, α2) eq → β
    = fun f x Refl → f x
```

With GADTs, the equality assumption is present at the term level, marked by a λ -abstraction over the type (α, α') `eq`, but the *use* of this equality is *implicit*: equality assumptions introduced by abstraction or pattern-matching can be silently used in the corresponding term clauses. This implicitness can be explained away by translating source terms into an intermediate language, such as System FC [Vytiniotis and Jones, 2011] that marks uses of equality assumptions with explicit coercions – providing a treatment similar to logical assumptions in proof assistants. But it can also be formalized directly, as in the formalization of GADTs extended to arbitrary logic constraints by Simonet and Pottier [2007], or Dependent ML by Xi [2007].

It is well-known however that, with implicit use of potentially-absurd assumptions, it is no longer safe to use full reduction under those assumptions. Assuming $(fst : (\alpha * \beta) \rightarrow \alpha)$ and $(true : bool)$, the term `apply fst true` reduces to

```
fun (Refl : (bool, (α * β)) eq) → fst true
```

Reducing under this abstraction would mean computing `fst true`, *i.e.* the application of a destructor to a constructor of an incompatible type, which is called a *runtime error*. Interestingly, this issue does not happen with an *explicit* handling of logical assumptions. In System FC, the above example would reduce to the following normal form (assuming that the assumption γ has been used to convert the type of the argument `true` rather than the type of the function `fst`):

```
fun (Refl (γ : bool ~# (α * β)) → fst (true ▷ γ))
```

Here, $bool \sim\# ('a * 'b)$ is the type of coercions that prove the equality between `bool` and $('a * 'b)$, and $(true \triangleright \gamma)$ is the application of the coercion γ to `true`. This application cannot be reduced until the formal variable γ has been instantiated (that is, never, if we are in an empty context with a sound type system). Meaning, γ remains in between `fst` and `true` preventing the application. Although System FC is a weak calculus (abstracting on term or coercion variables blocks reduction), full reduction could be used in a similar, explicit system.

We are convinced that it is important to also study the implicit presentation directly. There is a convergence of designs that indicates that *implicit* use of assumptions is significantly more convenient to the programmer. For a less-obvious example than GADTs in `ML` or `Haskell`, the book *Programming in Martin-Löf*

Type Theory [Nordström et al., 1990] uses a type theory with *extensional* equality, which allows implicit use of equality assumptions, especially to simplify programming with quotient types. We want to study λ -calculi that match how users wish to program and define the operational behavior of programs directly at this level.

Besides, we deem unfortunate the absolute reign of *weak reduction* on λ -calculi designed for programming. We argue that while one could have a weak-reduction semantics for reasoning about runtime complexity, a more abstract *full reduction* understanding is better to reason about correctness – as an important step towards equational reasoning on open terms.

In fact, type systems of programming languages are designed for full reduction strategies, and left unchanged when restricting the semantics to weak reduction strategies. For example, the type systems of ML, System F and its derivatives (F_η Mitchell [1988], $F_{<}$ Cardelli [1993], MLF Le Botlan and Rémy [2003], ...) are all sound for full reduction. While type systems are regularly improved to accept more well-typed programs, they do not try in general to take advantage of weak reduction strategies to accept nonsense under yet unapplied abstractions, *e.g.* $\lambda(x) 1$ true, on the basis that these errors won't be reachable by a weak-reduction strategy².

We claim that (pure) lambda-calculi for programming languages ought to strive to support *full reduction*; this design pressure should result in a better understanding of programming constructs. For example, soundness of full reduction subsumes soundness for any evaluation strategy such as call-by-name and call-by-value; and full reduction is used in practice in dependently-typed languages such as Coq or Agda, with significant efforts spent to make it practical Grégoire and Leroy [2002].

We could summarize the topic of this article with the following question. We know how to design calculi with *explicit* uses of logical assumptions and *full* reduction, or calculi with *implicit* uses of assumptions and *weak* reduction. Can we merge those apparently incompatible feature pairs into a single calculus, close to the non-encumbered terms the programmer wishes to write?

Consistent and inconsistent abstraction Intuitively, an abstraction on a type is *consistent* when we can prove at the point of abstraction that there always exists a possible instantiation for it in the current typing context; otherwise, we say it is *inconsistent*. A typical example of a consistent abstraction is an abstraction over a type variable α that has the kind \star of concrete types, as we know that at least `int` has kind \star so it is a valid instance for α . Abstraction over a type variable may also be constrained by a proposition that restricts the possible instances of the type variable. An example of an unsatisfiable proposition is the inter-convertibility $\text{int} \simeq (\tau \rightarrow \text{int})$, which is *absurd* for any type τ . Hence, an abstraction over a type variable α such that $\text{int} \simeq (\alpha \rightarrow \text{int})$ is inconsistent.

² One interesting counterexample is typechecking record concatenation, which delays the resolution of typing constraints based on the evaluation strategy Pottier [2000] in order to avoid the heavy cost of early consistency checking.

Previous work by Cretin [2014] and Cretin and Rémy [2014] introduced the calculus F_{cc} , built around *consistent coercion abstraction*, a mechanism that allows implicit abstraction over coercions and uses typing-transforming coercions, provided we prove at their abstraction point that they are instantiable; such coercions are completely erasable – they’re not at all present at the term level. This generalizes the traditional ML-style polymorphism in an expressive way, encompassing the type systems of System F, MLF, $F_{<:}$, F_η , and F_ω . To be able to abstract over hypotheses that may not be consistent, Cretin and Rémy also added a distinct mechanism of *inconsistent polymorphism* that is present at the term level and blocks reduction.

If an F_{cc} term a has type τ in the context $\Gamma, \alpha : \kappa$, and we can prove that the kind κ is instantiable by producing some type $\Gamma \vdash \sigma : \kappa$, we will consider that a has type $\forall(\alpha : \kappa) \tau$ in context Γ —the Curry-style presentation with no abstraction marker at the term level highlights that this form of polymorphism is erasable.

If on the contrary we do not know how to prove that κ is instantiable (or do not wish to), we may use *inconsistent* abstraction by building the term ∂a at the distinct type³ $\Pi(\alpha : \kappa) \tau$. This form blocks the reduction of a and is thus explicitly marked at the term level.

Discharging an inconsistent abstraction $\Pi(\alpha : \kappa) \tau$ also needs to be marked at the term level to unblock computation: if the type σ has kind κ , then κ is in fact inhabited and $a \diamond$ has type $\tau[\sigma/\alpha]$.

Even in full reduction (when reduction under λ ’s is allowed), reduction remains forbidden under ∂ ’s; an inconsistent abstraction is eliminated by the corresponding application, $(\partial a) \diamond$, which reduces to a letting the evaluation of a be resumed. In contrast, consistent abstraction is erasable by construction: it is absent from the term itself, which alone determines reduction.

Issues with F_{cc} In the absence of inconsistency abstraction, the language F_{cc} has a full reduction semantics and is confluent. However, both properties break when introducing inconsistent abstraction, since inconsistent abstraction blocks the evaluation to maintain soundness. This amounts to introducing a form of weak reduction inside the language. While some reductions under ∂ are unsound and must be blocked, others may be harmless and could be safely reduced—but this is not allowed. This is going against our claim that core calculi ought to support full reduction to the largest possible extent.

Besides, it is well known that mixing weak and strong reductions may break confluence, and this problem affects F_{cc} . If b reduces to b' , then $(\lambda(x) \partial x) b$ can reduce to either $(\lambda(x) \partial x) b'$ and then $\partial b'$, or to ∂b , which cannot be further reduced and, in particular, does not reduce to $\partial b'$, as confluence would require.

These issues were well-understood by Cretin and Rémy [2014] and left for future work. We present an improved variant of F_{cc} that solves both problems simultaneously. In the course of doing so, we also encountered some more minor

³ The following notation has nothing to do with dependent types.

issues in the details of F_{cc} , which allowed us to also improve the system as a whole.

Propositional truths and hiding The language F_{cc} uses a blocking construct ∂a to introduce the inconsistent abstraction $\Pi(\alpha : \kappa)\tau$. This does not match, however, the way potentially absurd assumptions are handled in dependent type theories, such as Coq, where reduction is blocked at the point of *use* of the assumption, not its point of *introduction*. This distinction is essential, in particular, to allow to write certified programs as Coq program terms: if the axioms (eg. classical logic or proof irrelevance) are only used in logic parts of the formalization (under terms at type **Prop**), they get removed by extraction; a program whose correctness proof uses axioms can compute – while it would be blocked if we used our ∂ to introduce the axiom.

We therefore split inconsistent abstraction into two more atomic notions. First, an abstraction form allows to *introduce* the assumption, but not yet to *use* it; this does not block computation – the assumption is just *frozen*. Second, an *elimination* construct on frozen assumptions makes them available for implicit use – blocking reduction.

Since the elimination construct blocks reduction, it needs to be present at the term level; it refers to assumption names introduced by the abstraction construct, which therefore also needs to be in terms – but without blocking reduction. In fact, we just reuse λ -abstraction for that purpose: locked assumptions are term variables at a new type $[P]$ of *propositional truths*, representing the assumption that the proposition P is true.

We write \diamond for the introduction of propositional truths, and $\delta(a, \phi.b)$ for its elimination. Informally,⁴ if the proposition P holds in the typing context Γ , then \diamond , at type $[P]$, is a witness of P . The corresponding elimination rule, $\delta(a, \phi.b)$ computes a at type $[P]$, while blocking the reduction of b , type-checked under the assumption $\phi : P$, until a turns into a concrete witness \diamond . Then, $\delta(\diamond, \phi.b)$ can be reduced to the pseudo-substitution $b[\diamond/\phi]$ whose effect is to remove all occurrences of ϕ in b , and finally the reduction can proceed.

With these new constructions, we may use standard abstraction $\lambda(x : [P]) a$ to abstract over a proposition P without blocking the evaluation of a , which means that a cannot use P yet. In particular, a may be of the form $a[\delta(x, \phi.b_1), \delta(x, \phi.b_2)]$, allowing the *implicit* use of the proposition P in subterms b_1 and b_2 , which cannot be reduced, while full reduction is still allowed in a .

Propositional truth elimination allows to express the fact that an assumption P may not actually be used directly at its abstraction site, but only “at some later time”. Conversely, there are situations where an elimination on P is needed to type-check parts of a term a and is no longer needed to typecheck some subterm b of a . To enable reduction of b , we introduce *assumption hiding* **hide ϕ in b** , which enforces that the proposition variable ϕ will *not* be used implicitly in the subterm b . In exchange for losing this convenience, we regain the full reduction behavior.

⁴ The language is formally defined in §2.

While assumption hiding has been introduced for programming reasons, it is also instrumental in restoring confluence. The loss of confluence happens when a substitution places a reducible term in an irreducible context. We may now restore confluence by inserting appropriate hidings during substitution when traversing proposition eliminators so as to preserve reducibility.

Contributions The central, novel idea of our work is the interaction of the explicit and implicit modes of use of logical assertions in a programming calculus admitting full-reduction. From a theoretical point of view, implicitness was a somewhat-neglected design choice, and we propose a continuum between implicit and explicit uses thanks to *propositional truths* and *assumption hiding* (Section 2.2). It reveals, for example, that GADTs are fundamentally different from the usual algebraic datatypes. From a practical point of view, this gives the user flexible control over the scope of logical assumptions to prevent them from leaking into unrelated parts of his program—while retaining the convenience of their implicit invocation.

Another, more technical contribution is a new formal full-reduction calculus F_{th} , with inconsistent coercion abstraction that is confluent (Section 3.5). It is notable that the construction that regains confluence (*hiding*) was initially motivated by programmer convenience.

Besides, there are several other contributions:

- We improve some details of the existing F_{cc} calculus, and were able to update its mechanized soundness proof accordingly (notes 5 and 6 in Section A.1). Although coercion calculi in the spirit of F_{cc} are neither surface nor internal languages, they are good at exploring the design space; hence, even small improvements are valuable in the long term.
- We extend (3.5) the confluence proof technique of Takahashi [1995] so that it scales to larger calculi expressed in the Wright-Felleisen style, using reduction contexts to factor out common patterns and avoid a combinatorial increase in the number of cases.
- When translating between two given calculi, precisely establishing a bisimulation generally requires the use of an administrative variant of the target calculus; in our case, we need an administrative arrow type that is incompatible with the usual arrow type. While this is a common trick in the literature, its soundness proof is not as obvious as one would expect. We provide precise proofs (A.2) that would be applicable to any calculus with several computational type constructors, *e.g.* arrows and products.

2 A calculus with propositional truths

In this section, we formally present our calculus, F_{th} – with propositional *truths* and *hiding*. As another instance of calculus based on erasable coercions, it is strongly inspired by the previous work on F_{cc} Cretin and Rémy [2014] and fol-

$\frac{\text{TERMVAR}}{\Gamma, x : \sigma, \Delta \vdash x : \sigma}$	$\frac{\text{TERMLAM}}{\Gamma \vdash \tau : \star \quad \Gamma, x : \tau \vdash a : \sigma \quad \Gamma \vdash \lambda(x) a : \tau \rightarrow \sigma}$	$\frac{\text{TERMAPP}}{\Gamma \vdash a : \tau \rightarrow \sigma \quad \Gamma \vdash b : \tau \quad \Gamma \vdash a b : \sigma}$
$\frac{\text{TERMPROD}}{\Gamma \vdash a : \tau_1 \quad \Gamma \vdash b : \tau_2 \quad \Gamma \vdash (a, b) : \tau_1 * \tau_2}$	$\frac{\text{TERMPROJ}}{\Gamma \vdash a : \tau_1 * \tau_2 \quad \Gamma \vdash \pi_i a : \tau_i}$	$\frac{\text{TERMCOERCE}}{\Gamma, \Sigma \vdash a : \tau \quad \Gamma \vdash (\Sigma \vdash \tau) \triangleright \sigma \quad \Gamma \vdash a : \sigma}$
$\frac{\text{TERMWIT}}{\Gamma \vdash Q \quad \Gamma \vdash \diamond : [Q]}$	$\frac{\text{TERMASSUME}}{\Gamma \vdash a : [P] \quad \Gamma, \phi : P \vdash b : \sigma \quad \Gamma \vdash \delta(a, \phi.b) : \sigma}$	$\frac{\text{TERMHIDE}}{\Gamma \Vdash \Delta \quad \Gamma \vdash \exists \Delta \quad \Gamma, \Delta \vdash a : \tau \quad \Gamma, \phi : P, \Delta \vdash \text{hide } \phi \text{ in } a : \tau}$

Fig. 1. F_{th} term typing judgment $\Gamma \vdash a : \tau$

$a, b ::= x, y \dots \mid \lambda(x) a \mid a a \mid (a, a) \mid \pi_i a$	Terms
$\mid \diamond \mid \delta(a, \phi.a) \mid \text{hide } \phi \text{ in } a$	
$\tau, \sigma ::= \alpha, \beta \dots \mid \tau \rightarrow \tau \mid \tau * \tau$	Types
$\mid \forall(\alpha : \kappa) \tau \mid (\tau, \sigma) \mid \pi_i \tau \mid () \mid [P]$	
$\kappa ::= \star \mid 1 \mid \kappa * \kappa \mid \{\alpha : \kappa \mid P\}$	Kinds
$P, Q ::= \top \mid P \wedge P \mid \forall(\alpha : \kappa) P \mid \exists \kappa \mid (\Sigma \vdash \tau) \triangleright \tau$	Prop.
$\Gamma, \Sigma, \Delta ::= \emptyset \mid \Gamma, x : \tau \mid \Gamma, \alpha : \kappa \mid \Gamma, \phi : P$	Contexts

Fig. 2. Syntax of terms, types, kinds and propositions

lows the same global structure of judgments. Yet, we do not assume familiarity with F_{cc} .⁵

We first present the general structure of judgments and the constructs that are common to both F_{cc} and F_{th} , together with their typing rules (§2.1). We then detail the novel features of F_{th} , namely propositional truths and assumption hiding (§2.2). Last, we present the dynamic semantics of F_{th} (§2.3). In §2.4, we introduce a variant of F_{th} that is used to prove the soundness of F_{th} by translation to F_{cc} in several steps (§3).

2.1 Consistent coercion calculus

Cretin and Rémy [2014] use a general notion of *erasable coercions* with abstraction over consistent coercions to present different type system features, such as polymorphism, subtyping, and more in a common framework where these features can be easily composed together. The restriction that only *consistent* coercions can be abstracted over is key to *erasability*.

Our calculus has four syntactic categories: terms a, b ; types τ, σ ; kinds κ ; and propositions P, Q . The syntax of each category and that of typing environments Γ , are described in Figure 2.

⁵ F_{cc} also supports equi-recursive types; we left them out of this presentation as they are orthogonal to reduction under inconsistent assumptions. It is the only feature of F_{cc} as previously described that is absent here.

The static semantics is given by four main judgments: a typing judgment $\Gamma \vdash a : \sigma$; a kinding judgment $\Gamma \vdash \sigma : \kappa$; a proposition satisfiability judgment $\Gamma \vdash P$; a coercion judgment $\Gamma \vdash (\Sigma \vdash \tau) \triangleright \sigma$; plus a context consistency judgment $\Gamma \vdash \exists \Delta$ and well-formedness judgments $\Gamma \Vdash t$ where t may be an environment, a kind, a proposition, or a coercion.

Terms We first describe terms of the consistent subset of F_{th} , which are the terms of the untyped λ -calculus with products, extended with one additional construct for coercions. Other constructs for manipulating inconsistent assumptions, namely, propositional truth and assumption hiding will be presented in §2.2.

The term typing judgment is defined by the rules in Figure 1. The introduction and elimination rules for arrows (TERMVAR, TERMLAM, TERMAPP) and products (TERMPROD, and TERMPROJ) are standard.

A remarkable feature of coercion calculi is that there is exactly one rule that does not change the term (and thus does not influence the dynamic semantics): the coercion rule TERMCOERCE. All runtime-irrelevant typing constructions, such as subtyping conversion and polymorphism introduction and elimination, are factorized into coercions. To express polymorphism, these coercions are *typing* coercions $(\Sigma \vdash \tau) \triangleright \sigma$ rather than *type* coercions $\tau \triangleright \sigma$: they also affect the typing environment Γ in which the coercion is used, by extending Γ with Σ when typechecking the premise of type τ , as described by Rule TERMCOERCE

This factorization has been explained in previous works of Cretin and Rémy [2014] and is orthogonal to our point of interest in the present paper, namely, the interplay between *program types* and *logical propositions* in a programming system. We thus focus our presentation on propositions in general rather than coercions, and propositional truths would naturally extend to many other program logics, such as arithmetic reasoning or general refinement types. Still, by maintaining a crisp separation between (Curry-style) *program terms* that compute and *derivations* on which we statically reason, consistent coercion calculi are good systems in which to think about implicit versus explicit uses of logic reasoning in program terms.

Coercions Despite the fact that coercions are included in the syntactic class of propositions, there are still two separate judgments.

The coercion judgment is defined in Figure 3. Besides structural rules of reflexivity (COERREFL) and transitivity (COERTRANS), coercions have rules for polymorphism (type abstraction COERGEN and type application COERINST), and for distributivity of coercions under *computational* type constructors (those that describe the shape of terms and appear in the term typing judgment): COERARROW, COERPROD, and COERWIT). Formulating those two orthogonal aspects as coercions allows to compose them easily: the F_η rules for instantiation of polymorphism under constructors naturally fall out as derived rules in consistent coercion calculi. Finally, Rule COERPROP allows to inject any propositional proof of a coercion (seen as a proposition) into the coercion judgment – when the coercion context is consistent. We refer the reader to Cretin and Rémy [2014] for detailed presentation of these rules.

$$\begin{array}{c}
\text{COERREFL} \\
\frac{\Gamma \vdash \tau \triangleright \tau}{\Gamma \vdash \tau \triangleright \tau} \\
\\
\text{COERINST} \\
\frac{\Gamma \vdash \sigma : \kappa}{\Gamma \vdash \forall(\alpha : \kappa) \tau \triangleright \tau[\sigma/\alpha]} \\
\\
\text{COERPROD} \\
\frac{\Gamma \vdash (\Sigma \vdash \tau) \triangleright \tau' \quad \Gamma \vdash (\Sigma \vdash \sigma) \triangleright \sigma'}{\Gamma \vdash (\Sigma \vdash \tau * \sigma) \triangleright \tau' * \sigma'} \\
\\
\text{COERTRANS} \\
\frac{\Gamma, \Sigma_1 \vdash (\Sigma_2 \vdash \tau_3) \triangleright \tau_2 \quad \Gamma \vdash (\Sigma_1 \vdash \tau_2) \triangleright \tau_1}{\Gamma \vdash (\Sigma_1, \Sigma_2 \vdash \tau_3) \triangleright \tau_1} \\
\\
\text{COERGEN} \\
\frac{\Gamma \vdash \exists \kappa}{\Gamma \vdash (\alpha : \kappa \vdash \tau) \triangleright \forall(\alpha : \kappa) \tau} \\
\\
\text{COERARROW} \\
\frac{\Gamma, \Sigma \vdash \tau' \triangleright \tau \quad \Gamma \vdash \tau : \star \quad \Gamma \vdash (\Sigma \vdash \sigma) \triangleright \sigma'}{\Gamma \vdash (\Sigma \vdash (\tau \rightarrow \sigma)) \triangleright (\tau' \rightarrow \sigma')} \\
\\
\text{COERWIT} \\
\frac{\Gamma, \phi : P \vdash Q}{\Gamma \vdash [P] \triangleright [Q]} \\
\\
\text{COERPROP} \\
\frac{\Gamma \vdash (\Sigma \vdash \tau) \triangleright \sigma \quad \Gamma \vdash \exists \Sigma}{\Gamma \vdash (\Sigma \vdash \tau) \triangleright \sigma}
\end{array}$$

Fig. 3. Coercion judgment $\Gamma \vdash (\Sigma \vdash \tau) \triangleright \sigma$

$$\begin{array}{c}
\text{PROPVAR} \\
\frac{\Gamma, \phi : P, \Delta \vdash P}{\Gamma \vdash P} \\
\\
\text{PROPAND} \\
\frac{\Gamma \vdash P_1 \quad \Gamma \vdash P_2}{\Gamma \vdash P_1 \wedge P_2} \\
\\
\text{PROP PROJ} \\
\frac{\Gamma \vdash P_1 \wedge P_2}{\Gamma \vdash P_i} \\
\\
\text{PROP GEN} \\
\frac{\Gamma \Vdash \kappa \quad \Gamma, \alpha : \kappa \vdash P}{\Gamma \vdash \forall(\alpha : \kappa) P} \\
\\
\text{PROP INST} \\
\frac{\Gamma \vdash \forall(\alpha : \kappa) P \quad \Gamma \vdash \tau : \kappa}{\Gamma \vdash P[\tau/\alpha]} \\
\\
\text{PROP TRUE} \\
\Gamma \vdash \top \\
\\
\text{PROP CONV} \\
\frac{\Gamma \vdash P \quad P =_{\beta} P' \quad \Gamma \Vdash P'}{\Gamma \vdash P'} \\
\\
\text{PROP KIND} \\
\frac{\Gamma \vdash \tau : \{\alpha : \kappa \mid P\}}{\Gamma \vdash P} \\
\\
\text{PROP INH} \\
\frac{\Gamma \vdash \sigma : \kappa}{\Gamma \vdash \exists \kappa} \\
\\
\text{PROP COER} \\
\frac{\Gamma \vdash (\Sigma \vdash \tau) \triangleright \sigma \quad \Gamma \vdash \tau : \star}{\Gamma \vdash (\Sigma \vdash \tau) \triangleright \sigma}
\end{array}$$

Fig. 4. Proposition satisfiability judgment $\Gamma \vdash P$

Notice how the introduction rule for polymorphism `COERGEN` requires the quantified-over kind κ to be inhabited—the proposition $\exists \kappa$ denoting kind inhabitation. This is the cornerstone of the distinction between *consistent* and *inconsistent* polymorphism: to abstract over a potentially-absurd kind or proposition (you have no inhabitation proof at hand), one must instead use the inconsistent abstraction, which changes the term as it blocks the reduction.

Kinding, satisfiability, and consistency Figures 5, 4, and 6 present those three related judgments.

The proposition satisfiability judgment $\Gamma \vdash P$ is defined in Figure 4. Besides coercions, the propositional features inherited from F_{cc} are relatively limited: there are the features used to subsume existing System F variants with some form of constrained quantification (F_{η} , $F_{<}$, $MLF \dots$), but more propositions could be added. The trivial true proposition, conjunction of propositions, and type-polymorphic propositions have obvious introduction and elimination rules.

The kind inhabitation proposition $\exists \kappa$ is true whenever kind κ is inhabited by some type σ ; we use the judgment $\Gamma \vdash \exists \kappa$ instead of $\Gamma \vdash \sigma : \kappa$ when only

$$\begin{array}{c}
\text{KINDVAR} \\
\frac{}{\Gamma, \alpha : \kappa, \Delta \vdash \alpha : \kappa} \\
\\
\text{KINDARROW} \\
\frac{\Gamma \vdash \tau : \star \quad \Gamma \vdash \sigma : \star}{\Gamma \vdash \tau \rightarrow \sigma : \star} \\
\\
\text{KINDPROD} \\
\frac{\Gamma \vdash \tau : \star \quad \Gamma \vdash \sigma : \star}{\Gamma \vdash \tau * \sigma : \star} \\
\\
\text{KINDWIT} \\
\frac{\Gamma \Vdash P}{\Gamma \vdash [P] : \star} \\
\\
\text{KINDALL} \\
\frac{\Gamma \Vdash \kappa \quad \Gamma, \alpha : \kappa \vdash \tau : \star}{\Gamma \vdash \forall(\alpha : \kappa) \tau : \star} \\
\\
\text{KINDUNIT} \\
\frac{}{\Gamma \vdash () : 1} \\
\\
\text{KINDPAIR} \\
\frac{\Gamma \vdash \tau : \kappa_1 \quad \Gamma \vdash \sigma : \kappa_2}{\Gamma \vdash (\tau, \sigma) : \kappa_1 * \kappa_2} \\
\\
\text{KINDPROJ} \\
\frac{}{\Gamma \vdash \pi_i \tau : \kappa_i} \\
\\
\text{KIND-REFINE} \\
\frac{\Gamma \vdash \tau : \kappa \quad \Gamma \Vdash \kappa \quad \Gamma, \alpha : \kappa \vdash P}{\Gamma \vdash \tau : \{\alpha : \kappa \mid P\}} \\
\\
\text{KIND-FORGET} \\
\frac{}{\Gamma \vdash \tau : \{\alpha : \kappa \mid P\}} \\
\\
\text{KINDCONV} \\
\frac{\kappa =_{\beta} \kappa' \quad \Gamma \vdash \tau : \kappa \quad \Gamma \Vdash \kappa'}{\Gamma \vdash \tau : \kappa'}
\end{array}$$

Fig. 5. Kinding judgment $\Gamma \vdash \tau : \kappa$

$$\begin{array}{c}
\text{CONTEEMPTY} \\
\frac{}{\Gamma \vdash \exists \emptyset} \\
\\
\text{CONTTERM} \\
\frac{\Gamma \vdash \exists \Delta \quad \Gamma \vdash \tau : \star}{\Gamma \vdash \exists(\Delta, x : \tau)} \\
\\
\text{CONTTYPE} \\
\frac{\Gamma \vdash \exists \Delta \quad \Gamma \vdash \exists \kappa}{\Gamma \vdash \exists(\Delta, \alpha : \kappa)} \\
\\
\text{CONTPROP} \\
\frac{\Gamma \vdash \exists \Delta \quad \Gamma \vdash P}{\Gamma \vdash \exists(\Delta, \phi : P)}
\end{array}$$

Fig. 6. Context consistency judgment $\Gamma \vdash \exists \Delta$

consistency matters. It is defined in Figure 4. Inhabitation is lifted to whole contexts in Figure 6: $\Gamma \vdash \exists \Delta$.

Kinding rules are defined in Figure 5. Kinding rules for base types are standard. The unit kind 1 is inhabited by the type-level trivial value $()$. Refinement kinds are the only construction introducing propositions in kinds—and thus in types: a refinement kind $\{\alpha : \kappa \mid P\}$ is inhabited by the types τ of kind κ such that the proposition $P[\sigma/\alpha]$ holds. For example, the bounded quantification $\forall(\alpha \leq \tau) \sigma$ can be expressed as $\forall(\alpha : \{\alpha : \star \mid \alpha \geq \tau\}) \sigma$. Refinement kinds also allow abstracting over propositions: assuming P in τ is expressed by refining the unit kind, $\Pi(\alpha : \{\alpha : 1 \mid P\}) \tau$.

Product kinds allow to quantify over several kinds at once. For example, the type $\forall(\alpha_1 : \kappa_1) \forall(\alpha_2 : \kappa_2) \tau$ can be seen as $\forall(\alpha : \kappa_1 * \kappa_2) \tau[(\pi_1 \alpha)/\alpha_1, (\pi_2 \alpha)/\alpha_2]$ but enforcing the two variables α_1 and α_2 to be introduced (and eliminated) simultaneously, which more often consistent than when the two variables are introduced one after the other.

The reader may have recognized in refinement kinds a restricted form of (kind-level) dependent product. Indeed, this would exactly be a dependent product if the propositions were included into the kinds – dependent products would then unify product kinds, refinement kinds. and conjunction of propositions. But F_{cc} 's irrelevant handling of proposition proofs allows for very simple, clutter-free elimination rules for the refinement kind, which do not have to appear in the syntax of types. We occasionally benefit from that convenience, for example to define the splitting operation (§A.1).

$$\begin{array}{c}
\Gamma \Vdash \star \quad \Gamma \Vdash 1 \quad \frac{\Gamma \Vdash \kappa_1 \quad \Gamma \Vdash \kappa_2}{\Gamma \Vdash \kappa_1 * \kappa_2} \quad \frac{\Gamma \Vdash \kappa \quad \Gamma, \alpha : \kappa \Vdash P}{\Gamma \Vdash \{\alpha : \kappa \mid P\}} \quad \Gamma \Vdash \top \\
\\
\frac{\Gamma \Vdash \kappa}{\Gamma \Vdash \exists \kappa} \quad \frac{\Gamma \Vdash P \quad \Gamma \Vdash Q}{\Gamma \Vdash P \wedge Q} \quad \frac{\Gamma \Vdash \kappa \quad \Gamma, \alpha : \kappa \Vdash P}{\Gamma \Vdash \forall (\alpha : \kappa) P} \\
\\
\frac{\Gamma \Vdash \Sigma \quad \Gamma, \Sigma \vdash \tau : \star \quad \Gamma \vdash \sigma : \star}{\Gamma \Vdash (\Sigma \vdash \tau) \triangleright \sigma} \quad \Gamma \Vdash \emptyset \quad \frac{\Gamma \Vdash \Delta \quad \Gamma, \Delta \vdash \tau : \star \quad x \notin \Gamma, \Delta}{\Gamma \Vdash \Delta, x : \tau} \\
\\
\frac{\Gamma \Vdash \Delta \quad \Gamma, \Delta \Vdash \kappa \quad \alpha \notin \Gamma, \Delta}{\Gamma \Vdash \Delta, \alpha : \kappa} \quad \frac{\Gamma \Vdash \Delta \quad \Gamma \Vdash P \quad \phi \notin \Gamma, \Delta}{\Gamma \Vdash \Delta, \phi : P}
\end{array}$$

Fig. 7. Well-formedness judgments

The presence of type-level data structures (in our case product kinds) implies a need for type-level computation and identification of computationally-equal objects, in particular in rules `PROP_CONV` and `KIND_CONV`. The conversion rules for kinds and propositions allow to convert between well-formed objects equal upto β -reduction of projections: $\pi_i(\tau_1, \tau_2) =_{\beta} \tau_i$, which is closed by congruence and equivalence to all types, propositions, and kinds.

Well-formedness Figure 7 presents the well-formedness judgments of F_{cc} for contexts, kinds and propositions, which are standard.

2.2 Propositional truths and hiding

The type $[P]$ represents the type of dynamic witnesses that P is satisfied. The type-checking rule for types of the form $[P]$ are listed in Figure 1. This type is introduced by the token \diamond , a ground value that inhabits $[P]$ exactly when the proposition P is satisfied in the current typing environment (Rule `TERM_WIT`). It is eliminated by the construction $\delta(a, \phi.b)$, where a must have a propositional truth type $[Q]$, and b is type-checked in an extended context where the assumption $\phi : Q$ is implicitly available – until it is hidden again in some subterm of the form `hide ϕ in ..`

As any other computational type, there is a distributivity coercion for propositional truths, Rule `COER_WIT` (Figure 3), which following F_{cc} design principle, can be derived and justified from the context $\delta(\square, \phi.\diamond)$, as an η -expansion of the identity context \square . The distributivity rule `COER_WIT` tells us that a witness for P , of type $[P]$, can be coerced into a witness for Q , of type $[Q]$, whenever P implies Q as a proposition.

Finally, the typing rule `TERM_HIDE` (Figure 1) for `hide ϕ in a` in context $\Gamma, \phi : P, \Delta$ is a form of weakening of ϕ . It is only valid under the condition that $\Gamma \vdash \exists \Delta$ holds. This does not mean that Δ must be consistent (it can depend on variables in Γ that were introduced by blocking elimination), but that it is consistent relative to Γ .

$$\begin{array}{l}
(\lambda(x) a) b \circ \rightarrow a[b/x]_{\emptyset} \\
\delta(\diamond, \phi.b) \circ \rightarrow b[\diamond/\phi] \\
\pi_i(a_1, a_2) \circ \rightarrow a_i \\
\\
c_{th} ::= \lambda(x) a \mid (a, b) \mid \diamond \\
d_{th} ::= \square b \mid \pi_i \square \mid \delta(\square, \phi.b) \\
\mathcal{E}_{th} \triangleq \{E[a] \mid \text{unguarded}(E), a = d_{th}[c_{th}], a \not\rightarrow\}
\end{array}
\quad
\begin{array}{l}
\text{CONTEXT} \\
\frac{a \circ \rightarrow b \quad \text{unguarded}(E)}{E[a] \longrightarrow E[b]}
\end{array}$$

Fig. 8. Dynamic semantics of \mathbb{F}_{th}

$$\begin{array}{ll}
\text{unguarded}(E) \triangleq (\text{guard}_{\emptyset}(E) = \emptyset) & \text{guard}_S(\delta(E, \phi.b)) \triangleq \text{guard}_S(E) \\
\text{guard}_S(\lambda(x) E) \triangleq \text{guard}_S(E) & \text{guard}_S(\delta(a, \phi.E)) \triangleq \text{guard}_{S \cup \{\phi\}}(E) \\
\text{guard}_S(a E) \triangleq \text{guard}_S(E) & \text{guard}_S(\text{hide } \phi \text{ in } E) \triangleq \text{guard}_{S \setminus \{\phi\}}(E) \\
\text{guard}_S(E a) \triangleq \text{guard}_S(E) & \text{guard}_S(\square) \triangleq S
\end{array}$$

Fig. 9. Guards

Kind-level propositions. *Propositional truths* are named as such because they are constructed and abstracted over in terms, with an explicit elimination construction; by contrast with the *definitional* judgment $\Gamma \vdash P$ which only lives in typing derivations. Note that it is possible to see propositions as kinds: the kind $\langle P \rangle$ defined as $\{\alpha : 1 \mid P\}$ is inhabited by $()$ exactly when the proposition P is satisfiable.

2.3 Dynamic semantics

The dynamic semantics of \mathbb{F}_{th} is defined in figures 8, 9 and 10. Because of assumption hiding, the notion of elimination contexts is non-standard: irreducible terms may have reducible subterms. In fact, the head β -reduction steps are also non-standard, because of the way hiding constructions are added during substitution of reducible values, so as to preserve confluence.

Reduction and head reduction We define a β -reduction relation (\longrightarrow) that is congruent to reduction contexts, and a *head* β -reduction relation ($\circ \rightarrow$) that only applies to head β -redexes (Figure 8). Distinguishing head reductions is important for the confluence proof (Section 3.5). Those reductions are fairly standard, except for the use of non-standard notions of substitution, and a side-condition on contexts described below.

Reduction contexts Full reduction is meant to allow any reduction path, so in general all one-hole term contexts E are reduction contexts. In \mathbb{F}_{th} , subterms that are in the scope of an implicit assumption, or equivalently of a proposition variable, must still be blocked. We use an auxiliary function $\text{guard}_S(E)$ to compute the set of proposition variables, called the *guards*, under which the hole \square of the

$$\begin{array}{ll}
x[b/x]_S \triangleq \mathbf{hide} S \mathbf{in} b & x[\diamond/\phi] \triangleq x \\
y[b/x]_S \triangleq y \quad (\text{if } y \neq x) & (\lambda(x) a)[\diamond/\phi] \triangleq \lambda(x) a[\diamond/\phi] \\
\diamond[b/x]_S \triangleq \diamond & (\mathbf{hide} \phi \mathbf{in} a)[\diamond/\phi] \triangleq a \\
(\lambda(y) a)[b/x]_S \triangleq \lambda(y) a[b/x]_S & (\mathbf{hide} \psi \mathbf{in} a)[\diamond/\phi] \triangleq \mathbf{hide} \psi \mathbf{in} a[\diamond/\phi] \\
(a a')[b/x]_S \triangleq a[b/x]_S a'[b/x]_S & \quad (\text{if } \psi \neq \phi) \\
\delta(a, \phi. a')[b/x]_S \triangleq \delta(a[b/x]_S, \phi. a'[b/x]_{S \cup \{\phi\}}) & \\
(\mathbf{hide} \phi \mathbf{in} a)[b/x]_S \triangleq \mathbf{hide} \phi \mathbf{in} a[b/x]_{S \setminus \{\phi\}} &
\end{array}$$

Fig. 10. Hiding and unhiding substitutions

single-hole context E is blocked, extended with an initial set S . The predicate $\mathbf{unguarded}(E)$ is then an abbreviation for the emptiness of $\mathbf{guard}_\emptyset(E)$. *Reduction contexts* are the *unguarded* one-hole contexts E . For example, $\delta(a, \phi.\square)$ is not a reduction context, whereas $(\lambda(x)\square)$ or $\delta(w_1, \phi.\delta(w_2, \psi.\mathbf{hide} \psi \mathbf{in} \mathbf{hide} \phi \mathbf{in} \square))$ are. Unguardedness is checked by an additional premise in Rule CONTEXT.

Hiding substitution $a[b/x]_S$ In order to preserve confluence it is essential that β -reduction preserves reducibility of subterms. A counter-example for confluence in F_{cc} , translated in F_{th} , is the term $(\lambda(x)\delta(y, \phi.x)) b$. The problem is that b appears in a reducible position but would become irreducible after one head reduction step, *i.e.* in the term $\delta(y, \phi.b)$ —with the usual notion of reduction.

Our solution is to define the reduction of λ -redexes using a non-standard notion of substitution, $a[b/x]_\emptyset$ that inserts assumption hidings as necessary for substituted terms to remain reducible. For instance, $\delta(y, \phi.x)[b/x]_\emptyset$ is equal to $\delta(y, \phi.\mathbf{hide} \phi \mathbf{in} b)$. In general, this hiding substitution can be indexed by any guard, which is the list of logical assumptions made so far during term traversal.

The hiding substitution is defined on Figure 10 where \triangleq stands for definition equality. Some cases that are simple traversals have been omitted.

Un-hiding substitution $b[\diamond/\phi]$ When the witness a of a propositional elimination $\delta(a, \phi.b)$ is blocked over reduces to \diamond , we know that the proposition witnessed by a is true, and the reduction of b can proceed. We remove each occurrence of $\mathbf{hide} \phi$ in each subterm of b , as it is not only unnecessary, but could also block now-reducible β -redexes if it remained: $a[\diamond/\phi]$ removes all occurrences of “ $\mathbf{hide} \phi$ ” in the term b while traversing b . It is also defined in Figure 10 – we give only a few representative cases. Note that the typing rule for assumption hiding guarantees that ϕ cannot appear in the subterm of $\mathbf{hide} \phi$ – and this property is preserved by reduction.

Errors Head reduction occurs when the destructor of some computational type meets the constructor of the same type. An *immediate error* is a term whose head is a destructor applied on a constructor of a different type. Figure 8 defines destructor contexts d_{th} and constructor terms c_{th} ; the set \mathcal{E}_{th} of *errors* is then defined as immediate errors occurring in a reduction context. Note that being

stuck on a free variable is not an error, that errors may still further reduce (if it contains other reducible positions with valid redexes), and that a non-error term may contain an immediate error blocked under a propositional elimination, such as $\lambda(x) \delta(x, \phi.\pi_1 \text{true})$ in our introductory example of abstracting over an equality between `int` and `bool`.

Below, we define errors for variants of our calculus in the same way, generated from a definition of constructor terms, destructor contexts, and the head reduction relation. Given a language of terms with a reduction relation and a set of errors \mathcal{E} , we say that a term a is *sound* if no reduction sequence starting from a ends in \mathcal{E} .

2.4 Two variants of \mathbf{F}_{th} : \mathbf{F}_t and \mathbf{F}_{cc}

The soundness of \mathbf{F}_{th} is proved by translation into \mathbf{F}_{cc} , which has been proved sound Cretin and Rémy [2014]. In fact, the translation is in two steps, using an intermediate calculus \mathbf{F}_t . Below, we formally define the calculi \mathbf{F}_t and \mathbf{F}_{cc} .

Removing assumption hiding. The language \mathbf{F}_t is obtain from \mathbf{F}_{th} by restricting to terms without hiding and by modifying the semantics of β -reduction so that it does not introduce hiding: the λ -reduction rule is $(\lambda(x) a) b \circ \rightarrow a[b/x]$. (As a consequence, \mathbf{F}_t is not confluent.)

The rest of the definition is unchanged; in absence of hiding, unguarded contexts $\mathbf{unguarded}(E)$ degenerate to a simpler, context-free definition that includes $\lambda(x)\square$ and $\delta(\square, \phi.b)$, but not $\delta(a, \phi.\square)$; and unhiding substitutions $b[\diamond/\phi]$ leave terms unchanged. Error terms \mathcal{E}_t are the subset of \mathcal{E}_{th} without hiding.

Primitive inconsistent abstraction. \mathbf{F}_{cc} uses a different primitive of *inconsistent abstraction* to work with inconsistent propositions, or rather potentially-uninhabited kinds. Its construction ∂a , mentioned in the introduction, blocks reduction immediately and has a type of the form $\Pi(\alpha : \kappa) \tau$ stating that it assumes a type $\alpha : \kappa$ while κ may be uninhabited. Conversely, $a \diamond$ unblocks a computation of type $\Pi(\alpha : \kappa) \tau$, whenever the kind κ can be shown inhabited. The head reduction rule is $(\partial a) \diamond \circ \rightarrow a$, and reduction contexts are as before, if we consider that each ∂E is guarded by a fresh propositional variable. The typing rules for those constructs are as follows:

$$\begin{array}{c}
\text{INCOHINTRO} \\
\frac{\Gamma \Vdash \kappa \quad \Gamma, \alpha : \kappa \vdash a : \tau}{\Gamma \vdash \partial a : \Pi(\alpha : \kappa) \tau} \\
\\
\text{INCOHELIM} \\
\frac{\Gamma \vdash a : \Pi(\alpha : \kappa) \tau \quad \Gamma \vdash \sigma : \kappa}{\Gamma \vdash a \diamond : \tau[\sigma/\alpha]} \\
\\
\text{KINDINCOH} \\
\frac{\Gamma \Vdash \kappa \quad \Gamma, \alpha : \kappa \vdash \tau : \star}{\Gamma \vdash \Pi(\alpha : \kappa) \tau : \star} \\
\\
\text{COERINCOH} \\
\frac{\Gamma, \alpha : \kappa', \Sigma \vdash \sigma : \kappa \quad \Gamma \vdash \exists \Sigma \quad \Gamma, \alpha : \kappa' \vdash (\Sigma \vdash \tau[\sigma/\alpha]) \triangleright \tau'}{\Gamma \vdash (\Sigma \vdash \Pi(\alpha : \kappa) \tau) \triangleright \Pi(\alpha : \kappa') \tau'}
\end{array}$$

The set \mathcal{E}_{cc} of \mathbf{F}_{cc} error terms is generated from its head reduction, its constructor terms (as in \mathbf{F}_{th} , but without \diamond and with (∂a)) and its destructor contexts (as in \mathbf{F}_{th} , but without $\delta(\square, \phi.b)$ and with $(\square \diamond)$).

3 Soundness and confluence

In this section we prove our two technical results, confluence and soundness of F_{th} . The proof proceeds by a series of translations, proving that the source language is sound if the target language is sound as well. In §3.1, we recall the F_{cc} soundness result from previous work. In §3.2, we show a translation from the sublanguage F_t to (an administrative variant of) F_{cc} . In §A.1, we show an inverse translation from F_{cc} to (an administrative variant of) F_t . That is, while F_t propositional truths are more convenient to use than F_{cc} 's inconsistent abstraction, they have equivalent behaviors. In §A.2, we show that the administrative variants of F_{cc} and F_t are sound if their non-administrative counterparts are. This establishes soundness of F_t . In §3.5, we prove confluence of F_{th} using parallel reductions. For this we precisely define F_{th} multi-hole contexts, which give convenient tools to reason on its dynamic semantics. Finally, §3.6 proves soundness of F_{th} , using the tools introduced for the confluence proof.

3.1 Soundness of F_{cc}

F_{cc} comes with a (computer-checked) soundness proof for its (non-deterministic) reduction: starting from a well-typed term, no reduction path can lead to an erroneous stuck term. For deep reasons detailed in the previous work, subject reduction (preservation of typing by reduction) does not hold for F_{cc} . Therefore, the soundness proof uses more semantic tools, building a model of the type system where types are sets of terms.

Theorem 1 (Previous work of Cretin and Rémy [2014]). *Soundness of F_{cc} Terms that are well-typed in F_{cc} in a consistent environment are sound. That is, if $\emptyset \vdash \exists \Gamma$ and $\Gamma \vdash a : \tau$, then a is sound.*

3.2 Translating propositional truths to F_{cc}

We now define a translation $\llbracket _ \rrbracket$ of terms, types and judgment derivations into F_{cc} . Informally, the idea of the translation is a form of CPS-encoding: we can translate a witness of type $[P]$ into a continuation consuming any inconsistent abstraction $\Pi(\alpha : \langle P \rangle) \tau$ to return a τ . Witness construction \diamond would become the elimination continuation $\lambda(x) (x \diamond)$, while propositional elimination $\delta(a, \phi.b)$ uses the translation of a as a continuation: $\llbracket a \rrbracket (\partial \llbracket b \rrbracket)$.

The actual translation on terms and types is close to the informal description above, with an important difference. The informal translation gives the expected computational behavior to well-typed terms, but has the defect of mapping some terms which are errors in F_t to terms in F_{cc} that may still further reduce: for example, $\delta((\lambda(x) x), \phi.y)$ is a stuck F_t term, but its translation $\llbracket \delta((\lambda(x) x), \phi.y) \rrbracket$, i.e. $(\lambda(x) x) (\partial y)$ can reduce.

Because the soundness proof of F_{cc} is done semantically, and subject reduction does not hold for this calculus, it is important that our translation of F_t terms be well-behaved even on ill-typed terms. Indeed, we will want to translate whole

reduction paths starting from a known F_t term which, even if well-typed, may reduce to ill-typed terms (but, as we prove in this section, not an error). We also want to reason about the translation of those sound but ill-typed reducts.

To get a translation of $\delta((\lambda(x) x), \phi.y)$ that is stuck, we use a slight variant of F_{cc} , called F_{cc}^b , for the target language. It is equipped with an “administrative” copy of the arrow type ($\tau \rightarrow^b \sigma$), of λ -abstraction ($\lambda^b(x) a$) and application ($a^b b$). The type system and reduction semantics are exactly those of F_{cc} , with each rule (in the static and dynamic semantics) about λ -abstractions duplicated into an identical “administrative” variant.

The administrative λ^b is entirely separate from the usual λ , and in particular $(\lambda(x) a)^b b$ and $(\lambda^b(x) a) b$ do not reduce and thus are both errors.

We can now formally define the translation from F_t to F_{cc}^b , which makes judicious use of administrative constructions to preserve stuck terms. It is defined below on the F_t -specific constructions; it just preserves the structure of other constructions and translate their subterms (we use $_$ for unused variable bindings):

$$\begin{aligned} \llbracket [P] \rrbracket &\triangleq \forall(\beta : \star) (\Pi(- : \{- : 1 \mid \llbracket [P] \rrbracket\}) \beta) \rightarrow^b \beta \\ \llbracket \delta(a, \phi.b) \rrbracket &\triangleq \llbracket [a] \rrbracket^b (\partial \llbracket [b] \rrbracket) \\ \llbracket [\diamond] \rrbracket &\triangleq \lambda^b(x) (x \diamond) \\ \llbracket [\Gamma, \phi : P] \rrbracket &\triangleq \llbracket [\Gamma] \rrbracket, \alpha : \{- : 1 \mid \llbracket [P] \rrbracket\} \end{aligned}$$

For example, the translation of the error $\delta((\lambda(x) x), \phi.y)$ is now $(\lambda(x) x)^b (\partial y)$, which is also an error. One cannot build a counter-example of the form $\delta((\lambda^b(x) a), \phi.b)$ as the administrative variants are not part of the input language F_t . One can show by induction that the translation preserves errors and typing.

Lemma 1 (Error preservation of F_t). *a is an error in F_t if and only if $\llbracket [a] \rrbracket$ is an error in F_{cc}^b .*

Lemma 2 (Typing preservation of F_t). *If $\Gamma \vdash a : \tau$ in F_t , then $\llbracket [\Gamma] \rrbracket \vdash \llbracket [a] \rrbracket : \llbracket [\tau] \rrbracket$ in F_{cc}^b .*

Proof. The translation $\llbracket [-] \rrbracket$ is extended to derivations in Figure 11. For each inference rule that proves a judgment from some premises, we show that the translation of the judgment is admissible from the translation of the premises. Direct induction shows that the translation of a valid derivation for the F_t judgment $\Gamma \vdash a : \tau$ is a valid derivation for the F_{cc}^b judgment $\llbracket [\Gamma] \rrbracket \vdash \llbracket [a] \rrbracket : \llbracket [\tau] \rrbracket$ – and mutually for all judgment forms.

Only the translations for inference rules of F_t -specific features are shown: the other inference steps are directly mapped in the expected way, for example

$$\left[\frac{\Gamma \vdash \tau_1 : \star \quad \Gamma \vdash \tau_2 : \star}{\Gamma \vdash \tau_1 * \tau_2 : \star} \right] \triangleq \frac{\llbracket [\Gamma] \rrbracket \vdash \llbracket [\tau_1] \rrbracket : \star \quad \llbracket [\Gamma] \rrbracket \vdash \llbracket [\tau_2] \rrbracket : \star}{\llbracket [\Gamma] \rrbracket \vdash \llbracket [\tau_1] * [\tau_2] \rrbracket : \star}$$

□

$$\begin{array}{c}
\left[\frac{\Gamma \Vdash P}{\Gamma \vdash [P] : \star} \right] \triangleq \frac{\frac{\frac{\frac{\llbracket \Gamma \rrbracket, \beta : \star, - : 1 \Vdash [P]}{\llbracket \Gamma \rrbracket, \beta : \star \vdash \Pi(- : \{- : 1 \mid [P]\}) \beta : \star}}{\llbracket \Gamma \rrbracket, \beta : \star \vdash \beta : \star}}{\llbracket \Gamma \rrbracket, \beta : \star \vdash (\Pi(- : \{- : 1 \mid [P]\}) \beta)} \rightarrow^b \beta : \star}}{\llbracket \Gamma \rrbracket \vdash \forall(\beta : \star) \Pi(\alpha : \{- : 1 \mid [P]\}) \rightarrow^b \beta : \star}} \\
\left[\frac{\Gamma \vdash a : [P] \quad \Gamma, \phi : P \vdash b : \sigma}{\delta(a, \phi.b) : \sigma} \right] \triangleq \\
\frac{\frac{\llbracket \Gamma \rrbracket \vdash [a] : \forall(\beta : \star) (\Pi(\alpha : \{- : 1 \mid [P]\}) \beta) \rightarrow^b \beta}{\llbracket \Gamma \rrbracket \vdash [a] : (\Pi(\alpha : \{- : 1 \mid [P]\}) [\sigma]) \rightarrow^b [\sigma]} \quad \frac{\llbracket \Gamma \rrbracket, \alpha : \{- : 1 \mid [P]\} \vdash [b] : [\sigma]}{\llbracket \Gamma \rrbracket \vdash \partial [b] : \Pi(\alpha : \{- : 1 \mid [P]\}) [\sigma]}}{\llbracket \Gamma \rrbracket \vdash [a]^b (\partial [b]) : [\sigma]} \\
\left[\frac{\Gamma \vdash P}{\Gamma \vdash \diamond : [P]} \right] \triangleq \frac{\frac{\frac{\frac{\llbracket \Gamma \rrbracket, \beta : \star, x : \Pi(\alpha : \{- : 1 \mid [P]\}) \beta \vdash [P]}{\llbracket \Gamma \rrbracket, \beta : \star, x : \Pi(\alpha : \{- : 1 \mid [P]\}) \beta \vdash 1 : \{- : 1 \mid [P]\}}}{\llbracket \Gamma \rrbracket, \beta : \star, x : \Pi(\alpha : \{- : 1 \mid [P]\}) \beta \vdash x \diamond : \beta}}{\llbracket \Gamma \rrbracket, \beta : \star \vdash \lambda^b(x) (x \diamond) : (\Pi(\alpha : \{- : 1 \mid [P]\}) \beta) \rightarrow^b \beta}}{\llbracket \Gamma \rrbracket \vdash \lambda^b(x) (x \diamond) : \forall(\beta : \star) (\Pi(\alpha : \{- : 1 \mid [P]\}) \beta) \rightarrow^b \beta}} \\
\left[\frac{\Gamma, \phi : P \vdash Q}{\Gamma \vdash [P] \triangleright [Q]} \right] \triangleq \\
\frac{\frac{\frac{\Gamma, \beta : \star \Vdash P \quad \Gamma, \beta : \star, - : \{- : 1 \mid [P]\} \vdash 1 : \{- : 1 \mid [Q]\} \quad \Gamma, \beta : \star, - : \{- : 1 \mid [P]\} \vdash \beta \triangleright \beta}{\Gamma, \beta : \star \vdash (\Pi(- : \{- : 1 \mid [Q]\}) \beta) \triangleright (\Pi(- : \{- : 1 \mid [P]\}) \beta)}}{\Gamma, \beta : \star \vdash ((\Pi(- : \{- : 1 \mid [P]\}) \beta) \rightarrow^b \beta) \triangleright ((\Pi(- : \{- : 1 \mid [Q]\}) \beta) \rightarrow^b \beta)}}{\Gamma \vdash (\forall(\beta : \star) (\Pi(- : \{- : 1 \mid [P]\}) \beta) \rightarrow^b \beta) \triangleright (\forall(\beta : \star) (\Pi(- : \{- : 1 \mid [Q]\}) \beta) \rightarrow^b \beta)}
\end{array}$$

Fig. 11. Translation of F_t into F_{cc}^b

Independently of typing preservation, we also prove a bisimulation property between F_t and F_{cc}^b . It is not quite the case that any single-reduction step in F_t is turned into a single-reduction step of F_{cc}^b , because the reduction of the translation of $\delta(\diamond, \phi.b)$, that is $(\lambda^b(x) (x \diamond))^b (\partial [b])$, does an extra administrative λ^b -reduction step before the expected ∂ -reduction. We define the relation $(\circ \rightarrow_b)$ of *administrative* head β -reductions as the subset, in F_{cc}^b , of reductions in $(\circ \rightarrow)$ of the form $(\lambda^b(x) a)^b b \circ \rightarrow a[b/x]$, and (\rightarrow_b) , the administrative β -reductions, its closure under reduction contexts. For any relation (\mathcal{R}) , we furthermore define the relation $\mathcal{R}^?$ by $a \mathcal{R}^? b$ if and only if $(a \mathcal{R} b) \vee (a = b)$.

Lemma 3 (Bisimulation of F_t by F_{cc}^b). *For any $a \rightarrow a'$ in F_t , we have $\llbracket a \rrbracket \rightarrow_b^? b \rightarrow \llbracket a' \rrbracket$ for some b in F_{cc}^b .*

Conversely, if $\llbracket a \rrbracket \longrightarrow b$ in F_{cc}^b , then $a \longrightarrow a'$ for some a' such that either $b = \llbracket a' \rrbracket$ or $\llbracket a \rrbracket \longrightarrow_b b \longrightarrow \llbracket a' \rrbracket$.

Proof. By lemma ??, the translation preserves context decomposition. Let us first show that we can, without loss of generality, study head reduction steps only.

If $E[a] \longrightarrow E[a']$ because $a \circ \rightarrow a'$, then it suffices to show that $\llbracket a \rrbracket \circ \rightarrow_b^? b \circ \rightarrow \llbracket a' \rrbracket$ to get $\llbracket E[a] \rrbracket \circ \rightarrow_b^? \llbracket E \rrbracket [b] \circ \rightarrow \llbracket E[a'] \rrbracket$, as $\llbracket E \rrbracket$ is an reduction context.

In the other direction, we must be careful in presence of administrative reductions. If there is no administrative step and $\llbracket E[a] \rrbracket \longrightarrow \llbracket E \rrbracket [b]$ because $b = \llbracket a' \rrbracket$, it suffices to show $a \circ \rightarrow a'$. If there is an administrative step, $\llbracket E[a] \rrbracket \longrightarrow_b \llbracket E \rrbracket [b] \longrightarrow \llbracket b' \rrbracket$ with $a \circ \rightarrow_b b$, it is not immediate that b' is indeed of the form $\llbracket E[a'] \rrbracket$ for some $b \circ \rightarrow \llbracket a' \rrbracket$: the second reduction could happen in a different reducible position. But this cannot happen, because not reducing the ∂ -redex that results from the administrative reduction would give a term that is not in the image of the translation: all λ^b -redexes in translated terms come from $\llbracket \delta(\diamond, \phi.a') \rrbracket = (\lambda^b(x)(x \diamond))^b (\partial \llbracket a' \rrbracket)$, and reducing those always creates a ∂ -redex, which cannot happen in a translated term. So a single-reduction step of the form $\llbracket E \rrbracket [(\partial \llbracket a' \rrbracket) \diamond] \longrightarrow \llbracket b' \rrbracket$ is only possible if $\llbracket b' \rrbracket = \llbracket E[a'] \rrbracket$. We therefore have $\llbracket a \rrbracket \circ \rightarrow_b b \circ \rightarrow \llbracket a' \rrbracket$ and it suffices to show $a \circ \rightarrow a'$.

For redexes that are present in both F_t and F_{cc}^b , we can prove the forward and backward simulation at the same time. We show that the translation of any F_t -redex maps its F_t -reduction to a F_{cc}^b -reduction, and conversely that any F_{cc}^b -reduction on this redex, from a translated term, goes to a translated term and can be mapped back into a F_t -reduction. We use the substitution lemma ?? for reductions defined with substitutions:

- $\pi_i(a, b) \circ \rightarrow a_i$ and $\llbracket \pi_i(a, b) \rrbracket = \pi_i(\llbracket a \rrbracket, \llbracket b \rrbracket) \circ \rightarrow \llbracket a_i \rrbracket$; this is the only head $(,)$ -redex in the image of the F_t translation.
- $(\lambda(x)a) b \circ \rightarrow a[b/x]$ and $\llbracket (\lambda(x)a) b \rrbracket = (\lambda(x) \llbracket a \rrbracket) \llbracket b \rrbracket \circ \rightarrow \llbracket a \rrbracket \llbracket \llbracket b \rrbracket / x \rrbracket = \llbracket a[b/x] \rrbracket$; this is the only head λ -redex in the image of the translation.

(Note that, in the cases above, no head administrative reduction could start from the translation of the source of the head transition considered.)

In the case of \diamond -redexes $\delta(\diamond, \phi.b) \circ \rightarrow b$, we have

$$\begin{aligned} & \llbracket \delta(\diamond, \phi.b) \rrbracket \\ &= (\lambda^b(x)(x \diamond))^b (\partial \llbracket b \rrbracket) \\ &\circ \rightarrow_b (\partial \llbracket b \rrbracket) \diamond \\ &\circ \rightarrow \llbracket b \rrbracket \end{aligned}$$

Conversely, no head \diamond -redex may appear in a translated term, and the only λ^b -redex comes from the translation of a \diamond -redex, where it reduces to a ∂ -redex as in the reduction sentence above. □

Corollary 1. *If $\llbracket a \rrbracket$ is sound in F_{cc}^b , then a is sound in F_t .*

Proof. Suppose $a \longrightarrow^* b$, we have to show that $b \notin \mathcal{E}_{F_t}$. By forward simulation 3 we have that $\llbracket a \rrbracket \longrightarrow^* \llbracket b \rrbracket$; if $\llbracket a \rrbracket$ is sound in F_{cc}^b , then $\llbracket b \rrbracket \notin \mathcal{E}_{F_{cc}^b}$. By forward preservation of errors 1, we conclude that $b \notin \mathcal{E}_{F_t}$. \square

Corollary 2. *If F_{cc}^b is sound, then so is F_t .*

Proof. We must show that well-typed terms in F_t do not go wrong. If a is well-typed in F_t , then by preservation of typing (lemma 2) $\llbracket a \rrbracket$ is well-typed in F_{cc}^b . Assuming soundness of F_{cc}^b , $\llbracket a \rrbracket$ is thus sound. We conclude that a is sound by the previous corollary 1. \square

We note that we do not need the bisimulation result to establish soundness (relative to F_{cc}^b), but only the forward simulation and the forward translation of errors.

The backward simulation shows that besides having the same soundness property, F_t and F_{cc}^b are also the same in term of number of reductions up to administrative steps: reasoning on program efficiency can therefore also be transposed from one to the other. In this respect, it may be important to remark that the one computation step we allowed to neglect, the administrative λ^b -reduction, never performs arbitrary duplication of its argument: whenever it appears in the translation, the λ^b -variable appears exactly once in the body. We could better enforce this invariant by using a linear type for this administrative construction, but this would require invasive changes to the type system.

3.3 Translating F_{cc} into F_t

Just as we presented a translation from F_t into (an administrative variant of) F_{cc} to prove F_t 's soundness, it is possible and enlightening to translate F_{cc} back into (an administrative variant of) F_t – after fixing a minor defect of F_{cc} as previously presented. For lack of space, we have not included this translation in the present document, but it is available in the full version (and in Appendix §A.1).

3.4 Soundness of the administrative arrow

To conclude, from the two previous sections, that F_{cc} 's soundness implies F_t 's soundness and conversely, we need to prove the soundness of the administrative variants relative to their base calculus. While this is a common technique, this result should not be neglected, and its soundness proof is actually not as obvious as one would expect. By lack of space, the proof is only available in the long version of this document (and in Appendix §A.2).

This result proves, in particular, the soundness of F_{cc}^b relative to F_{cc} . Along with Corollary 2, establishing the soundness of F_t relative to F_{cc}^b , and the already established soundness of F_{cc} (Theorem 1) this concludes the soundness proof of F_t (Theorem 4, §4).

$$\begin{array}{c}
\Box_i : S \vdash \Box_i : S \quad \emptyset \vdash x : S \quad \frac{\Gamma \vdash E : S \setminus \{\phi\}}{\Gamma \vdash \mathbf{hide} \phi \mathbf{in} E : S} \quad \frac{\Gamma \vdash E_1 : S \quad \Delta \vdash E_2 : S \cup \{\phi\}}{\Gamma, \Delta \vdash \delta(E_1, \phi.E_2) : S} \\
\\
\frac{\Gamma \vdash E_1 : S \quad \Delta \vdash E_2 : S}{\Gamma, \Delta \vdash (E_1 E_2), (E_1, E_2) : S} \quad \frac{\Gamma \vdash E : S}{\Gamma \vdash (\lambda(x) E), (\pi_i E), (\sigma_i E) : S} \\
\\
\frac{a = E[x]^i \quad x \notin E \quad (\Box_i : S_i) \vdash E : S}{a[b/x]_S \triangleq E[\mathbf{hide} S_i \mathbf{in} b]^i} \quad \frac{a = E[\mathbf{hide} \phi \mathbf{in} b_i]^i \quad \phi \notin E}{a[\diamond/\phi] \triangleq E[b_i]^i}
\end{array}$$

Fig. 12. Guard analysis of multi-hole contexts

3.5 Confluence of \mathbf{F}_{th}

Multi-hole contexts Figure 12 introduces a new judgment $(\Box_i : S_i)^{i \in I} \vdash E : S$, that is a simple syntactic analysis of the *guards* of a multi-hole context, that is the set of propositional variables that block the reduction of each hole. The judgment can be read as “if the whole term is guarded by S , then the i -th hole \Box_i is guarded by S_i ”. A multi-hole context is just a term whose variables are, by convention, named \Box_i for some i in I , and which appear only once in the term; we enforce that latter invariant by using disjoint union for the context union Γ, Δ , which corresponds to a simple linear typing discipline. The notation $E[_]^{i \in I}$ corresponds to a context with a family of holes indexed by i , and in contexts $\amalg^{i \in I} \Delta_i$ is the disjoint union of a family of contexts $(\Delta_i)^{i \in I}$. For sake of brevity, we often leave I implicit and just write i instead of $i \in I$.

Notice that $\mathbf{guard}_S(E)$ for a single-hole context is uniquely defined by $\Box : \mathbf{guard}_S(E) \vdash E : S$. We also use multi-contexts to re-define the hiding substitution $a[b/x]_S$ defined in §2.3, and the hide-removing substitution $a[\diamond/\phi]$ used in the reduction rule for $\delta(\diamond, \phi.a)$.

Finally, we say that a multi-context E is a *prefix* of another multi-context E' (or a term, if E' has no holes) if E' can be obtained by substituting sub-contexts into the holes of E .

Parallel reductions We prove confluence using the Tait-Martin-Löf technique of parallel reductions, with a simple proof argument inspired by Takahashi [1995]. The idea of Takahashi is that parallel reduction (noted $a \Longrightarrow b$) for the simple λ -calculus with only arrows can be made deterministic by adding a redex-avoidance rule (the $a \neq \lambda(-)$ - hypothesis below, meaning that a does not start with an abstraction) to the parallel reduction of application:

$$\frac{a \neq \lambda(-) \quad a \Longrightarrow a' \quad b \Longrightarrow b'}{a b \Longrightarrow a' b'} \quad \frac{a \Longrightarrow a' \quad b \Longrightarrow b'}{(\lambda(x) a) b \Longrightarrow a' [b'/x]}$$

Without the redex-avoiding condition $a \neq \lambda(-)$ in the application reduction rule, two reduction paths are available to each β -redex, performing the β -reduction or not. This gives a parallel condition that *may* reduce each redex of the term in one step, and in particular subsumes the usual single-step reduction relation

$$\begin{array}{c}
\frac{E \not\rightarrow \quad (\square_i : \emptyset)^i \vdash E : \emptyset \quad (a_i \circ\Rightarrow b_i)^i \quad (a_i \Longrightarrow a'_i)^i \quad R_\beta[a'_i]^i \overset{?}{\circ\rightarrow} b}{E[a_i]^i \Longrightarrow E[b_i]^i \quad R_\beta[a_i]^i \circ\Rightarrow b} \\
R_\beta ::= (\lambda(x) \square_1) \square_2 \mid \pi_i(\square_1, \square_2) \mid \delta(\diamond, \phi.b)
\end{array}$$

Fig. 13. Parallel reduction

by choosing to reduce exactly one redex. Takahashi remarks that the condition forces all redexes of the term to be reduced (Gross-Knuth reduction), and that this modified relation trivially forces confluence of the parallel reduction of which it is a special case.

Adapting Takahashi's idea to a Wright-Felleisen setting of head reduction and elimination contexts suggests a new formulation, which is to decompose a reducible term a into the form $E[b_i]^i$ where the multi-context E is *not reducible* when seen as a term – a generalization of the redex-avoiding condition. For the same reason that Takahashi's reduction was deterministic, the decomposition of a into $E[b_i]^i$ where E is not reducible and the b_i are head redexes is unique, since E is the largest head context that does not contain redexes. This way to capture the set of redexes of a term lets us define parallel reduction rather than only the Gross-Knuth reduction.

Figure 13 gives the definition of our parallel reduction $a \Longrightarrow b$, mutually defined with the *head* parallel reduction $a \circ\Rightarrow b$ that reduces head redexes – and other redexes in the subterms. The notation $E \not\rightarrow$ can be understood in term of the single-step reduction relation, when E is seen as a term as any other: $\neg(\exists E', E \rightarrow E')$.

The parallel reduction of $E[a_i]^i$ only happens when the a_i are all redexes, as they must be related to some b_i by the head parallel reduction ($\circ\Rightarrow$) that only starts from head redexes $R_\beta[\square_i]^i$. Not all these redexes need to be reduced, however, as the head beta-reduction step $R_\beta[a'_i]^i \overset{?}{\circ\rightarrow} b$ is optional. In particular, taking $R_\beta[a'_i]^i = b$ for each redex shows that the relation (\Longrightarrow) is reflexive.

The restriction that the substituted terms a_i are redexes is crucial to modularly reason about reducibility; for if we substituted the non-redex $\lambda(x) a$ into the context $(\square b)$, we would get a reducible result while neither the term nor the context were. No such situation can happen when the plugged terms are head redexes themselves, as redexes do not overlap.

Lemma 4 (Orthogonality). *Redexes do not overlap: If $(\square_i : \emptyset)^i \vdash E : \emptyset$ is a one-hole irreducible context distinct from \square , then for any redex contexts $R[\square_j]^j$ and $(R'_i[\square_k]^k)^i$ and families of terms $(a_j)^j$ and $(b_{i,k})^{i,k}$ we have $R[a_j]^j \neq E[R'_i[b_{i,k}]^k]^i$.*

Proof. Immediate by case analysis of all pairs (R, E) of a redex context and an irreducible context. For example, if $R = (\lambda(x) \square_1) \square_2$, then:

- if E has $(\lambda(x) \square_1) \square_2$ as head prefix, it is reducible
- if E is not a prefix of $(\lambda(x) \square_1) \square_2$, filling its hole can never match R

– if E is $(\square_3 b)$ or $(\square_3 \square_4)$, no head redex put in \square_3 could match $(\lambda(x) \square_1)$

□

The other lemma we need to prove the unicity of the decomposition by irreducible contexts is about the structure of reducible positions in a term or context.

Lemma 5 (Reducible positions).

For any guard S , any term a has a minimal non-empty prefix F such that $(\square_k : \emptyset)^k \vdash F[\square_k]^k : S$. For any non-empty prefix F' of a with $(\square_{k'})^{k'} \vdash F'[\square_{k'}]^{k'} : S$, F is a prefix of F' . Furthermore, F is irreducible.

Proof. A direct induction shows that for any a and S , there exists a unique derivation of $\vdash a : S$.

Then, proceeding by induction on this derivation, it suffices to cut its smallest prefix F whose variables have guard \emptyset . This prefix is minimal by construction (the derivation $\vdash a : S$ being unique).

It is immediate that F is irreducible; for if it had a reducible subterm, F would be of the form $G[b]$ where b has a head-redex and is in reducible position in F . We would then have $\square : \emptyset \vdash G : S$ as a strict prefix of F , which is absurd.

Lemma 6 (Unique decomposition of irreducible contexts). *If two parallel reductions have the same source, then they use the same context-redexes decomposition.*

Proof. If we have

$$\frac{E \not\rightarrow (\square_i : \emptyset)^i \vdash E : \emptyset \quad (a_i \rightleftharpoons b_i)^i}{E[a_i]^i \Longrightarrow E[b_i]^i}$$

$$\frac{E' \not\rightarrow (\square'_j : \emptyset)^j \vdash E' : \emptyset \quad (a'_j \rightleftharpoons b'_j)^j}{E'[a'_j]^j \Longrightarrow E'[b'_j]^j}$$

with $E[a_i]^{i \in I} = E'[a'_j]^{j \in J}$, we want to prove that $I = J$, $E = E'$ and $\forall i, a_i = a'_i$.

The proof proceeds by induction, and is independent from the details of \mathbf{F}_{th} as its useful properties were expressed by the two previous lemmas. The base case of the induction uses orthogonality (Lemma 4) and the other cases are uniformly handled by a non-empty minimal context of reducible positions (Lemma 5).

If E is \square , the whole term has a head redex; we know by orthogonality 4 that E' must be \square as well, so the context decomposition is unique. The situation if E' is \square is symmetrical.

Otherwise, we know that $E \neq \square$ and $E' \neq \square$. We take the minimal non-empty prefix F (Lemma 5) of $E[a_i]^i = E'[a'_j]^j$ such that $(\square_k : \emptyset)^{k \in K} \vdash F : \emptyset$. As it is a minimal prefix, we have sub-families $(E_k)^k, (E'_k)^k$ and partitions $((a_{i_k})^{i_k \in I_k})^{k \in K}$ (with $I = \coprod^k I_k$) and $((a'_{j_k})^{j_k \in J_k})^{k \in K}$ (with $J = \coprod^k J_k$) of $(a_i)^i$ and $(a'_j)^j$ such that

$$F[E_k[a_{i_k}]^{i_k}]^k = E[a_i]^i = E'[a'_j]^j = F[E'_k[a'_{j_k}]^{j_k}]^k$$

In particular, for each $k \in K$ we have $E_k[a_{i_k}]^{i_k} = E'_k[a'_{j_k}]^{j_k}$. F is non-empty so the E_k, E'_k are strictly smaller than E, E' , so by induction hypothesis we get $I_k = J_k$, $E_k = E'_k$ and $(a_{i_k})^{i_k} = (a'_{j_k})^{j_k}$. This lets us conclude by

$$I = \Pi^k I_k = \Pi^k J_k = J$$

and

$$E = F[E_k]^k = F[E'_k]^k = E'$$

and

$$(a_i)^i = \Pi^k (a_{i_k})^{i_k} = \Pi^k (a'_{j_k})^{j_k} = (a'_j)^j$$

□

In the general case of filling a context E with subterms that are not necessarily head redexes, we may still reason on reducibility of the subterms:

Lemma 7 (Composability of parallel reduction).

The following rule, which does not constrain the $(a_i)^i$ to be head redexes or E to be irreducible, is admissible:

$$\frac{(\Box_i : \emptyset)^i \vdash E : \emptyset \quad (a_i \Longrightarrow b_i)^i}{E[a_i]^i \Longrightarrow E[b_i]^i}$$

Proof. Let the context/redex decomposition of each a_i be $F_i[a'_j]^{j \in J_i}$, such that we have for each i

$$\frac{(a'_j \circ \Rightarrow b'_j)^{j \in J_i} \quad (\Box_j : \emptyset)^j \vdash F_i : \emptyset \quad F_i \not\rightarrow}{a_i = F_i[a'_j]^j \Longrightarrow F_i[b'_j]^j = b_i}$$

then we can conclude by the following admissible derivation (Lemma ??):

$$\frac{(a'_j \circ \Rightarrow b'_j)^{i \in I, j \in J_i} \quad (\Box_j : \emptyset)^{i \in I, j \in J_i} \vdash E[F_i]^i : \emptyset}{E[a_i]^i \Longrightarrow E[b_i]^i}$$

□

The last technical lemma we need closes a commutative diagram between parallel reduction and one-step head reduction.

Lemma 8 (Commutation \Longrightarrow and $\circ \rightarrow$). *If $R_\beta[a_i]^i \Longrightarrow R_\beta[a'_i]^i$ and both $R_\beta[a_i]^i \circ \rightarrow b$ and $R_\beta[a'_i]^i \circ \rightarrow b'$, then $b \Longrightarrow b'$.*

Proof. $R_\beta[a_i]^i \circ \Rightarrow R_\beta[a'_i]^i$ exactly means that for any i we have $a_i \Longrightarrow a'_i$. We then reason by case analysis on R_β .

The result is immediate if R_β reduce to a \Box_i context; for example, if R_β is $\pi_2(\Box_1, \Box_2)$, we have $R_\beta[a_i]^i \circ \rightarrow a_2$ and $R_\beta[a'_i]^i \circ \rightarrow a'_2$, so our goal is $a_2 \Longrightarrow a'_2$ which is exactly our assumption.

If the head reduction involves a substitution, we need to prove that substitution preserves reducible positions, which is a non-trivial fact that crucially relies on the hiding construct. Consider the λ -reduction $(\lambda(x) a_1) a_2$: we need to prove that $a_1[a_2/x]_\emptyset \Longrightarrow a'_1[a'_2/x]_\emptyset$. This proceeds in two steps:

$$a_1[a_2/x]_\emptyset \Longrightarrow a_1[a'_2/x]_\emptyset \Longrightarrow a'_1[a'_2/x]_\emptyset$$

Let $E[\square_i]^i$ be the context of x 's occurrences in a_1 : $a_1 = E[x]^i$. There is a unique family $(S_i)^i$ of guards such that $(\square_i : S_i)^i \vdash E : \emptyset$, and the context in which a_2 reduces in $a_1[a_2/x]_\emptyset$ is, by construction, $E[\mathbf{hide} S_i \mathbf{in} \square'_i]^i$. By substitution of the guard judgment we have:

$$\frac{(\square_i : S_i)^i \vdash E : \emptyset \quad (\square'_i : \emptyset \vdash \mathbf{hide} S_i \mathbf{in} \square'_i : \emptyset)^i}{(\square'_i : \emptyset)^i \vdash E[\mathbf{hide} S_i \mathbf{in} \square'_i]^i : \emptyset}$$

so

$$a_1[a_2/x]_\emptyset = E[\mathbf{hide} S_i \mathbf{in} a_2]^i \Longrightarrow E[\mathbf{hide} S_i \mathbf{in} a'_2]^i = a_1[a'_2/x]_\emptyset$$

The second step, $a_1[a'_2/x]_\emptyset \Longrightarrow a'_1[a'_2/x]_\emptyset$, comes from the fact that substituting a variable with a term in a potentially-reducible context does not change its guards: if $(\square_i : S_i) \vdash E : S$, and a contains no \square_i , then $(\square_i : S_i) \vdash E[a/x] : S$ as well. This transformation can be applied to all contexts involved in the derivation of $a_1 \Longrightarrow a_2$; the substituted contexts may become reducible, so we use the composability lemma (7) to rebuild a derivation of $a_1[a'_2/x]_\emptyset \Longrightarrow a'_1[a'_2/x]_\emptyset$. \square

Note that it is precisely that last lemma that failed with F_t or F_{cc} without a hiding construct. Indeed, with $b \Longrightarrow b'$, reducing $(\lambda(x) \delta(y, \phi.x)) b$ to $\delta(y, \phi.b)$ does not allow to close the diagram to $\delta(y, \phi.b')$, while reducing to $\delta(y, \phi.\mathbf{hide} \phi \mathbf{in} b)$ allows to close by reducing to $\delta(y, \phi.\mathbf{hide} \phi \mathbf{in} b')$.

Theorem 2. *The parallel reduction relation (\Longrightarrow) is confluent.*

Proof. We prove, by mutual induction, that:

- (\Longrightarrow) is confluent: if $a \Longrightarrow b^1$ and $a \Longrightarrow b^2$, then there exists a b' such that $b^1 \Longrightarrow b'$ and $b^2 \Longrightarrow b'$
- (\Longrightarrow) is (\Longrightarrow)-confluent: if $a \rightrightarrows b^1$ and $a \rightrightarrows b^2$, then there exists a b' such that $b^1 \Longrightarrow b'$ and $b^2 \Longrightarrow b'$

In the case of (\Longrightarrow), by unicity of context/redex decomposition (Lemma 6), the two reductions are of the form

$$\frac{E \not\rightarrow (\square_i : \emptyset)^i \vdash E : \emptyset \quad (a_i \rightrightarrows b^1_i)^i}{E[a_i]^i \Longrightarrow E[b^1_i]^i} \quad \frac{E \not\rightarrow (\square_i : \emptyset)^i \vdash E : \emptyset \quad (a_i \rightrightarrows b^2_i)^i}{E[a_i]^i \Longrightarrow E[b^2_i]^i}$$

where only the head parallel reductions $(a_i \rightrightarrows b^1_i)^i$ and $(a_i \rightrightarrows b^2_i)^i$ may differ. By mutual induction hypothesis (\Longrightarrow -confluence of (\rightrightarrows), for each i there is a b'_i such that $b^1_i \Longrightarrow b'_i$ and $b^2_i \Longrightarrow b'_i$. We can then conclude by defining

$b' \triangleq E[b'_i]^i$, as we have $E[b_i^1]^i \Longrightarrow E[b'_i]^i$ and $E[b_i^2]^i \Longrightarrow E[b'_i]^i$ by composability of the parallel reduction (Lemma 7).

In the case of $(\circ\Rightarrow)$, let us write $a \overset{\epsilon}{\circ\rightarrow} b$, for $\epsilon \in \{0, 1\}$, when $a \overset{?}{\circ\rightarrow} b$, with $\epsilon = 1$ when there is a $(\circ\rightarrow)$ step and $\epsilon = 0$ when $a = b$.

We have two reductions

$$\frac{(a_i \Longrightarrow a_i^{\prime 1})^i \quad R_\beta[a_i^{\prime 1}]^i \overset{\epsilon^1}{\circ\rightarrow} b^1}{R_\beta[a_i]^i \circ\Rightarrow b^1} \quad \frac{(a_i \Longrightarrow a_i^{\prime 1})^i \quad R_\beta[a_i^{\prime 1}]^i \overset{\epsilon^2}{\circ\rightarrow} b^1}{R_\beta[a_i]^i \circ\Rightarrow b^1}$$

and want a b' such that $b^1 \circ\Rightarrow b'$ and $b^2 \circ\Rightarrow b'$. By mutual induction hypothesis (confluence of (\Longrightarrow)), we have a family $(a'_i)^i$ such that $a_i^{\prime 1} \Longrightarrow a'_i$ and $a_i^{\prime 2} \Longrightarrow a'_i$. Let us now define b' to be the unique term such that $R_\beta[a'_i]^i \overset{\epsilon^1 \vee \epsilon^2}{\circ\rightarrow} b'$ – we could simply define it such as $R_\beta[a'_i]^i \overset{1}{\circ\rightarrow} b'$, which gives the Gross-Knuth reduction used in Takahashi's proof; but handling the case where $\epsilon^1 \vee \epsilon^2 = 0$ is immediate, the important case is 1, so this wouldn't simplify that much.

We can check by case analysis on ϵ^1 and ϵ^2 that $b^1 \circ\Rightarrow b'$ and $b^2 \circ\Rightarrow b'$; we will only prove $b^1 \circ\Rightarrow b'$, the other result being symmetrical.

– If $\epsilon^1 = 0$, then $\epsilon^1 \vee \epsilon^2 = \epsilon^2$, and we have

$$\frac{(a_i^{\prime 1} \Longrightarrow a'_i)^i \quad R_\beta[a'_i]^i \overset{\epsilon^2}{\circ\rightarrow} b'}{b_1 = R_\beta[a_i^{\prime 1}]^i \circ\Rightarrow b'}$$

so in particular $b^1 \Longrightarrow b'$.

– If $\epsilon_1 = 1$, then we have $R_\beta[t_i^{\prime 1}]^i \Longrightarrow R_\beta[t'_i]^i \circ\Rightarrow b'$ and $R_\beta[t_i^{\prime 1}]^i \circ\rightarrow b^1$, so we can conclude $b^1 \Longrightarrow b'$ from the commutation lemma 8.

□

Corollary 3 (Confluence). *The relation (\longrightarrow^*) is confluent.*

Proof. We only need to prove that $(\Longrightarrow^*) = (\longrightarrow^*)$, as Pollack Pollack [1995] proved in a generic way that for any relation R , if R is confluent then R^* is confluent as well – so we know that (\Longrightarrow^*) is confluent.

The proof proceeds by double inclusion. We prove that $(\longrightarrow) \subset (\Longrightarrow)$, and (by induction) that $(\Longrightarrow) \subset (\longrightarrow^*)$.

If $a \longrightarrow b$, then $a = E[a']$ and $b = E[b']$ for some reduction context E with $a' \circ\rightarrow b'$, which immediately implies $a' \circ\rightarrow b'$ (no reduction of the subterms, and one $(\circ\rightarrow)$ -reduction). We can then conclude $a = E[a'] \Longrightarrow E[b'] = b$ by head reduction through an arbitrary context (Lemma ??).

In the other direction, we first remark that for any reduction multi-context $E[\square_i]^i$ and family of reductions $(a_i)^i \longrightarrow^* (b_i)^i$, one has $E[a_i]^i \longrightarrow^* E[b_i]^i$; indeed we have

$$\begin{aligned} E[a_i]^i &= E[a_1, a_2, \dots, a_n] \longrightarrow^* E[b_1, a_2, \dots, a_n] \\ &\longrightarrow^* E[b_1, b_2, \dots, a_n] \longrightarrow^* \dots \longrightarrow^* E[b_1, b_2, \dots, b_n] = E[b_i]^i \end{aligned}$$

Then we prove by mutual induction that both (\implies) and (\impliedby) are included in (\longrightarrow^*) . In a derivation of $R_\beta[a_i]^i \impliedby R[a'_i]^i \longrightarrow b$, we can inductively convert all the assumption $(a_i \implies a'_i)^i$ into $(a_i \longrightarrow^* a'_i)^i$, which gives us $R_\beta[a_i]^i \longrightarrow^* R_\beta[a'_i]^i \longrightarrow b$ from the previous remark. In a derivation of $E[a_i]^i \implies E[b_i]^i$ from $(a_i \impliedby a'_i)^i$, we get $(a_i \longrightarrow^* a'_i)^i$ by induction hypothesis, and our remark lets us conclude with $E[a_i]^i \longrightarrow^* E[b_i]^i$.

3.6 Soundness of \mathbf{F}_{th}

The soundness proof of \mathbf{F}_{th} is again a translation from \mathbf{F}_{th} to \mathbf{F}_t with a forward simulation. Before getting to the translation proper, we need to study two transformations used to define it. *Hide-extrusion* (3.6) removes hiding from a \mathbf{F}_{th} term, and its correctness property let us simulate forward reductions of the form $\delta(\diamond, \phi.b) \circ \rightarrow b[\diamond/\phi]$. *Hide-normalization* (3.6) strengthens the structure of hiding in a \mathbf{F}_{th} term, in such a way that we can forward-simulate the other \mathbf{F}_{th} reductions, despite the mismatch between \mathbf{F}_{th} 's hiding substitution $a[b/x]_S$ and \mathbf{F}_t 's natural substitution. We finally prove \mathbf{F}_{th} 's soundness (Theorem 3).

Hide-extrusion In a language without hiding such as \mathbf{F}_t , it is possible for the programmer to emulate the effects of hiding by extruding terms out of a blocking construction. Instead of $\delta(a, \phi.E[\mathbf{hide} \phi \mathbf{in} b])$, one can write $\mathbf{let} \ x_b = b \ \mathbf{in} \ \delta(a, \phi.E[x_b])$, where b appears in reducible position; we call this transformation *hide-extrusion*. In the general case, E may bind variables or block over other proposition variables, and the translation needs to be refined to preserve b 's typing environment; for example, $\delta(a, \phi.\lambda(y) (f (\mathbf{hide} \phi \mathbf{in} b)))$ is hide-extruded into $\mathbf{let} \ x_b = \lambda(y) b \ \mathbf{in} \ \delta(a, \phi.\lambda(y) (f (x_b y)))$.

Figure 14 gives a formal definition of the hide-extruding rewrite $a \hookrightarrow a'$ by defining two functions $\mathbf{abs}(C, a)$, which abstracts a over all the variables bound in the context C , and $\mathbf{app}(b, C)$, which closes such an abstracted b (when applied from under the context C) by applying it to the appropriate variables, accordingly. These definitions are factorized by a grammar N of context frames that do not bind any variable. If a has type τ , then $\mathbf{abs}(C, a)$ has type $\mathbf{arr}(C, \tau)$, and conversely if b has type $\mathbf{arr}(C, \tau)$ then $\mathbf{app}(b, C)$ has type τ .

Lemma 9 (Typing preservation of hide-extrusion).

If a is well-typed in \mathbf{F}_{th} and $a \hookrightarrow b$, then b is well-typed, at the same type.

Proof. This is easily established from the two following rules, shown derivable by a direction induction on C .

$$\frac{\Gamma, \square : \tau \vdash C[\square] : \sigma \quad \Gamma \vdash a : \tau}{\Gamma \vdash \mathbf{abs}(C, a) : \mathbf{arr}(C, \tau)} \quad \frac{\Gamma, \square : \tau \vdash C[\square] : \sigma \quad \Gamma \vdash b : \mathbf{arr}(C, \tau)}{\Gamma \vdash C[\mathbf{app}(b, C)] : \sigma}$$

□

Lemma 10 (Hide-extrusion of errors). *If $a \hookrightarrow b$, then a is an error if and only if b is an error.*

$$\begin{aligned}
& \delta(b, \phi.C[\mathbf{hide} \phi \mathbf{in} a]) \leftrightarrow \mathbf{let} \ x = \mathbf{abs}(C, a) \ \mathbf{in} \ \delta(b, \phi.C[\mathbf{app}(x, C)]) \\
N ::= & \ \square \ b \mid a \ \square \mid (\square, b) \mid (a, \square) \mid \pi_i \ \square \mid \sigma_i \ \square \mid \delta(\square, \phi.b) \quad \text{Non-binding contexts} \\
& \mathbf{abs}(\square, a) \triangleq a \\
& \mathbf{abs}(N[C], a) \triangleq \mathbf{abs}(C, a) \\
& \mathbf{abs}(\lambda(y) C, a) \triangleq \lambda(y) \mathbf{abs}(C, a) \\
& \mathbf{abs}(\delta(w, \psi.C), a) \triangleq \lambda(x_w) \delta(x_w, \psi.\mathbf{abs}(C, a)) \\
& \mathbf{arr}(\square, \tau) \triangleq \tau \\
& \mathbf{arr}(N[C], \tau) \triangleq \mathbf{arr}(C, \tau) \\
& \mathbf{arr}(\lambda(y : \sigma') C, \tau) \triangleq \sigma' \rightarrow \mathbf{arr}(C, \tau) \\
& \mathbf{arr}(b_1((y_1 : \sigma'_1).C \mid y_2.b_2), \tau) \triangleq \sigma'_1 \rightarrow \mathbf{arr}(C, \tau) \\
& \mathbf{arr}(x_{b_1}(y_1.b_1 \mid (y_2 : \sigma'_2).C), \tau) \triangleq \sigma'_2 \rightarrow \mathbf{arr}(C, \tau) \\
& \mathbf{arr}(\delta((w : [P]), \psi.C), \tau) \triangleq [P] \rightarrow \mathbf{arr}(C, \tau) \\
& \mathbf{app}(a, \square) \triangleq a \\
& \mathbf{app}(a, N[C]) \triangleq \mathbf{app}(a, C) \\
& \mathbf{app}(a, \lambda(y) C) \triangleq \mathbf{app}(a \ y, C) \\
& \mathbf{app}(a, \delta(x_w, \psi.C)) \triangleq \mathbf{app}(a \ \diamond, C)
\end{aligned}$$

Fig. 14. Hide-extrusion: translating **hide** back into plain F_t

Proof. The destructors in reducible positions are exactly the same in a term and its extrusion – $\mathbf{abs}(C, a,)$ only introduces new constructors on top of the existing subterm a . *A fortiori* the reducible error redexes are the same.

□

Lemma 11 (Extrusion Reduction).
 $\mathbf{app}(\mathbf{abs}(C, a), C) \longrightarrow^* a[\diamond/\mathbf{guard}_\emptyset(C)]$.

For all term a , we have

Proof. For convenience, let us define $S(C) \triangleq \mathbf{guard}_\emptyset(C)$. Then, by induction on C :

$$\begin{aligned}
& \mathbf{app}(\mathbf{abs}(F[C], a), F[C]) \\
= & \mathbf{app}(\mathbf{abs}(C, a), C) \\
\longrightarrow_{\text{hyp}} & a[\diamond/S(C)] \\
\\
& \mathbf{app}(\mathbf{abs}((\lambda(y) C), a), \lambda(y) C) \\
= & \mathbf{app}((\lambda(y) \mathbf{abs}(C, a)) y, C) \\
\longrightarrow & \mathbf{app}(\mathbf{abs}(C, a)[y/y]_\emptyset, C) \\
= & \mathbf{app}(\mathbf{abs}(C, a[y/y]_{\mathbf{guard}_\emptyset(\mathbf{abs}(C, \square))}), C) \\
= & \mathbf{app}(\mathbf{abs}(C, a[y/y]_{S(C)}), C) \\
= & \mathbf{app}(\mathbf{abs}(C, a[(\mathbf{hide} S(C) \mathbf{in} y)/y]), C) \\
\longrightarrow_{\text{hyp}} & a[(\mathbf{hide} S(C) \mathbf{in} y)/y][\diamond/S(C)] \\
= & a[\diamond/S(C)] \\
\\
& \mathbf{app}(\mathbf{abs}(\delta(w, \phi.C), a), \delta(w, \phi.C)) \\
= & (\lambda(x_w) \delta(x_w, \phi.\mathbf{abs}(C, a))) \diamond \\
\longrightarrow & \mathbf{app}(\delta(\diamond, \phi.\mathbf{abs}(C, a)), C) \\
\longrightarrow & \mathbf{app}(\mathbf{abs}(C, a[\diamond/\phi]), C) \\
\longrightarrow_{\text{hyp}} & a[\diamond/\phi][\diamond/S(C)] \\
= & a[\diamond/(\phi, S(C))]
\end{aligned}$$

In the $\lambda(y) C$ case, we used the fact that the only free occurrences of y in $\mathbf{abs}(C, b)$ are in b , which is checked by immediate induction. \square

Each hide-extrusion rewrite removes exactly one $\mathbf{hide} \phi$ from the source term; in particular, iterating hide-extrusion terminates, and gives a term without any hiding construct. We prove in 3.6 that this gives a forward simulation of \mathbf{F}_{th} by \mathbf{F}_t . This is easy to see in the simple case of extrusion through a reduction context E without any other $\mathbf{hide} \phi$:

$$\delta(b, \phi.E[\mathbf{hide} \phi \mathbf{in} a]) \hookrightarrow \mathbf{let} x = \mathbf{abs}(E, a) \mathbf{in} \delta(b, \phi.E[\mathbf{app}(x, E)])$$

The only reducible subterms of the source term are a and b . b is still reducible in the target term, and a as well as $\mathbf{guard}_\emptyset(\mathbf{abs}(E, \square)) = \mathbf{guard}_\emptyset(E) = \emptyset$. If b eventually reduces to \diamond , then the source becomes $E[a]$ which may reduce further. The target can reduce to $\mathbf{let} x = \mathbf{abs}(E, a) \mathbf{in} E[\mathbf{app}(x, E)]$, which itself reduces to $E[\mathbf{app}(\mathbf{abs}(E, a), E)]$, then to $E[a]$ by the previous lemma 11.

Hide-normalization The remaining issue for a forward simulation of \mathbf{F}_{th} by \mathbf{F}_t is the difference between the substitutions used in β -reductions. If $(\lambda(x) a) b$ is related to some $(\lambda(x) a') b'$ by hide-extrusion, $a[b/x]_\emptyset$ may not be related to $a'[b'/x]$ in the general case, as the substitution in \mathbf{F}_{th} may introduce new hiding constructs that have to be extruded again.

The idea of hide-normalization is to rewrite a term so that both substitutions *coincide*, by establishing the invariant that the guard of each bound variable occurrence is equal to the guard of its binder. For example, in $\lambda(x)\delta(y, \phi.x)$ the guard of x 's binding site is \emptyset , while its occurrence has guard $\{\phi\}$. β -reducing this λ -abstraction would introduce a `hide` ϕ . We can statically rewrite it into $\lambda(x)\delta(y, \phi.\mathbf{hide}\ \phi\ \mathbf{in}\ x)$, which is equivalent (unblocking free variables doesn't affect reduction), and whose β -reduction doesn't introduce hiding.

In the general case, we define the hide-normalization function $\mathbf{H}(a)$ from \mathbf{F}_{th} to \mathbf{F}_{th} . It recursively traverses all subterms and is a direct mapping, except in the λ -case where it is defined as:

$$\frac{\text{HIDENORMLAM} \quad (\square_i : S_i)^i \vdash C : \emptyset \quad x \notin C}{\mathbf{H}(\lambda(x) C[x]^i) \hat{=} \lambda(x) \mathbf{H}(C[\mathbf{hide}\ S_i\ \mathbf{in}\ x]^i)}$$

Lemma 12 (Type preservation of hide-normalization). *If a is well-typed in \mathbf{F}_{th} , then $\mathbf{H}(a)$ is also well-typed, at the same type.*

Proof. If $(\square_i : S_i)^i \vdash E : \emptyset$, we consider the derivation tree of

$$\frac{\Gamma, x : \tau \vdash E[x]^i : \sigma}{\Gamma \vdash \mathbf{H}(\lambda(x) E[x]^i) : \tau \rightarrow \sigma}$$

The i -th x -variable rule, occurring under the guard $S_i = (\phi_j)_{j \in J_i}$, is of the form

$$\Gamma, x : \tau, \Delta, (\phi_j, \Delta_j : P_j)^j \vdash x : \tau$$

where the context is ordered: Δ has no propositional variables, and the ϕ_j are the propositional variables introduced after x in the context, and the Δ_j the context fragments introduced before the next propositional variable. By this ordering, we know that well-formedness of x does not depend on the Δ_j , so we may strengthen this variable rule, without affecting its validity, into

$$\Gamma, x : \tau, \Delta, (\phi_j : P_j)^j \vdash x : \tau$$

To get a valid derivation of $\Gamma \vdash E[\mathbf{hide}\ S_i\ \mathbf{in}\ x]^i$, we simply replace all such leaves by a valid sub-derivation

$$\frac{\Gamma, x : \tau \vdash x : \tau}{\Gamma, x : \tau, (\phi_j : P_j)^j \vdash \mathbf{hide}\ S_i\ \mathbf{in}\ x : \tau}$$

□

Lemma 13 (Error preservation of hide-normalization). *A \mathbf{F}_{th} term a is an error if and only if $\mathbf{H}(a)$ is an error.*

Proof. This is immediate, as hide-normalization only affects bound variables. It does not change the structure of constructor or destructors, and thus preserves errors.

□

Lemma 14 (Hide-normalization is stable by reduction). *If $\mathbb{H}(a) \longrightarrow b'$, then b' is equal to $\mathbb{H}(b)$ for some b .*

Proof. The “well-hiding” judgment $\Gamma \Vdash a : S$

By construction, hide-normalized terms a are exactly those that satisfy $\Gamma \Vdash a : \emptyset$, for the following “well-hiding” judgment $\Gamma \Vdash a : S$:

$$\begin{array}{c}
\text{WHVAR} \\
\frac{}{\Gamma, x : S, \Delta \Vdash x : S} \\
\\
\text{WHLAM} \\
\frac{\Gamma, x : S \Vdash a : S}{\Gamma \Vdash \lambda(x) a : S} \\
\\
\text{WHAPP} \\
\frac{\Gamma \Vdash a : S \quad \Gamma \Vdash b : S}{\Gamma \Vdash a b : S} \\
\\
\text{WHPROD} \\
\frac{\Gamma \Vdash a : S \quad \Gamma \Vdash b : S}{\Gamma \Vdash (a, b) : S} \\
\\
\text{WHPROJ} \\
\frac{\Gamma \Vdash a : S}{\Gamma \Vdash \pi_i a : S} \\
\\
\text{WHWIT} \\
\frac{}{\Gamma \Vdash \diamond : S} \\
\\
\text{WHASSUME} \\
\frac{\Gamma \Vdash a : S \quad \Gamma, \phi \Vdash b : S \cup \{\phi\}}{\Gamma \Vdash \delta(a, \phi.b) : S} \\
\\
\text{WHHIDE} \\
\frac{\Gamma, \Delta \Vdash a : S \setminus \{\phi\}}{\Gamma, \phi, \Delta \Vdash \text{hide } \phi \text{ in } a : S}
\end{array}$$

Defining hide-normalization in terms of well-hiding

To prove properties of hide-normalization we will need a more precise definition. A natural definition is to recursively define $\mathbb{H}(\Gamma, a, S)$, where Γ is a context of the well-hiding judgment and S a guard, such that $\mathbb{H}(a)$ is $\mathbb{H}(\emptyset, a, \emptyset)$, with the following characteristic property

$$\forall \Gamma, a, S, \text{freevars}(a) \subset \Gamma \implies \Gamma \Vdash \mathbb{H}(\Gamma, a, S) : S$$

The definition is as follows:

$$\begin{aligned}
\mathbb{H}((\Gamma, x : S, \Delta), x, S') &\triangleq \text{hide } S' \setminus S \text{ in } x \\
\mathbb{H}(\Gamma, \lambda(x) a, S) &\triangleq \lambda(x) \mathbb{H}((\Gamma, x : S), a, S) \\
\mathbb{H}(\Gamma, a b, S) &\triangleq \mathbb{H}(\Gamma, a, S) \mathbb{H}(\Gamma, b, S) \\
\mathbb{H}(\Gamma, (a, b), S) &\triangleq (\mathbb{H}(\Gamma, a, S), \mathbb{H}(\Gamma, b, S)) \\
\mathbb{H}(\Gamma, \pi_i a, S) &\triangleq \pi_i \mathbb{H}(\Gamma, a, S) \\
\mathbb{H}(\Gamma, \diamond, S) &\triangleq \diamond \\
\mathbb{H}(\Gamma, \delta(a, \phi.b), S) &\triangleq \delta(\mathbb{H}(\Gamma, a, S), \phi. \mathbb{H}((\Gamma, \phi), b, S \cup \{\phi\})) \\
\mathbb{H}((\Gamma, \phi, \Delta), \text{hide } \phi \text{ in } a, S) &\triangleq \text{hide } \phi \text{ in } \mathbb{H}((\Gamma, \Delta), a, S \setminus \{\phi\}) \\
\mathbb{H}(\Gamma, \square, S) &\triangleq (\Gamma \Vdash S)
\end{aligned}$$

The last case does not make sense on terms, but it extends the $\mathbb{H}(\Gamma, -, S)$ on contexts: $\mathbb{H}(\Gamma, E[\square], S)$ is a context $E'[\Gamma' \Vdash S']$ such that, for any a , $\mathbb{H}(\Gamma, E[a], S) = E'[\mathbb{H}(\Gamma', a, S')]$.

Subject reduction of $(\emptyset \Vdash \emptyset)$ We now show subject reduction for the judgment $\emptyset \vdash a : \emptyset$, which shows that hide-normalization is stable by reduction.

It is easily checked by induction that unguarded single-hole contexts $\mathbf{unguarded}(E)$ are exactly those that satisfy

$$\emptyset \Vdash E[\Gamma \Vdash \emptyset] : \emptyset$$

Now suppose we have a reduction

$$\frac{a \circ \rightarrow b \quad \mathbf{unguarded}(E)}{E[a] \longrightarrow E[b]}$$

we have to show that, for any Γ such that $\emptyset \Vdash E[\Gamma \Vdash \emptyset] : \emptyset$, head reduction $a \circ \rightarrow b$ preserves well-guardedness under E , that is $\Gamma \Vdash a : \emptyset$ implies $\Gamma \Vdash b : \emptyset$. We do a case analysis on $a \circ \rightarrow b$, reasoning by inversion on the $\Gamma \Vdash a : \emptyset$ derivation.

This is immediate for $\pi_i(a_1, a_2) \circ \rightarrow a_i$: $\Gamma \Vdash \pi_i(a_1, a_2) : \emptyset$ has hypothesis $\Gamma \Vdash \pi_i(a_1, a_2) : \emptyset$, which itself has both hypotheses $\Gamma \Vdash a_1 : \emptyset$ and $\Gamma \Vdash a_2 : \emptyset$.

For $\delta(\diamond, \phi.b) \longrightarrow b[\diamond/\phi]$, inversion on $\Gamma \Vdash \delta(\diamond, \phi.b) \Vdash \emptyset$ gives us the hypothesis $\Gamma, \phi : \emptyset \Vdash b : \phi$. We can then inductively check that if $\Gamma, \phi \Vdash a : S \cup \{\phi\}$ holds, then $\Gamma \Vdash a[\diamond/x] : S$ holds as well. This gives us $\Gamma \Vdash b[\diamond/\phi] : \emptyset$ as expected.

For $(\lambda(x)a)b \circ \rightarrow a[b/x]_\emptyset$, inversion on $\Gamma \Vdash (\lambda(x)a)b : \emptyset$, gives us the hypotheses $\Gamma, x : \emptyset \Vdash a : \emptyset$ and $\Gamma \Vdash b : \emptyset$, and we need to prove that $\Gamma \Vdash a[b/x]_\emptyset : \emptyset$.

To prove this, we first check inductively that if we have $\Gamma, x : S_x, \Delta \Vdash E[x]^i : S$ with $x \notin E$, then the multi-hole guard judgment $(\square_i : S_x)^i \vdash E : S$ holds – all occurrences of the context’s hole are under the same guard, which is characteristic of hide-normalization.

In particular, from hypothesis $\Gamma, x : \emptyset \Vdash a : \emptyset$ we know that a is of the form $E[x]^i$ with $x \notin E$ and $(\square : \emptyset)^i \vdash E : \emptyset$. Using the characterization of hiding substitution in terms of multi-contexts 3.5, we then have $E[x]^i[b/x]_\emptyset \hat{=} E[b]^i$ (in each whole we put $\mathbf{hide} \emptyset \mathbf{in} b$, which is exactly b). In other words, for our hide-normalized a and b ,

$$a[b/x]_\emptyset = a[b/x]$$

This is the characteristic property of hide-normalized terms: because we have added all the necessary hiding on variables already, hiding substitution never needs to add more, and coincides with the usual substitution.

We then conclude by checking inductively that the well-hiding judgment is substitutive: for any S , if $\Gamma, x : S_x \Vdash a : S$ and $\Gamma \Vdash b : S_x$, then $\Gamma \Vdash \Gamma[b/x] : S$. \square

Lemma 15 (Hide-normalization is a forward simulation). *If $a \longrightarrow b$ then $\mathbf{H}(a) \longrightarrow \mathbf{H}(b)$.*

Proof. This proof uses the tools of the previous proof that hide-normalization is preserved by reduction 14.

We prove the following statement, strengthened for induction: for any Γ, S such that $\text{freevars}(a) \subset \Gamma$, if $a \longrightarrow b$ then $\mathbb{H}(\Gamma, a, S) \longrightarrow \mathbb{H}(\Gamma, b, S)$.

If we have a reduction

$$\frac{a \circ \rightarrow b \quad \text{unguarded}(E)}{E[a] \longrightarrow E[b]}$$

We know that $\mathbb{H}(\Gamma, E[\square], S)$ is of the form $E'[\Gamma' \vdash \emptyset]$, because E' is both normalized and unguarded and thus verifies $\emptyset \Vdash E'[\Gamma' \vdash \emptyset] : \emptyset$ for some Γ' . It then suffices to show that $\mathbb{H}(\Gamma', a, \emptyset) \circ \rightarrow \mathbb{H}(\Gamma', b, \emptyset)$, and we can conclude by

$$\mathbb{H}(\Gamma, E[a], \emptyset) = E'[\mathbb{H}(\Gamma', a, \emptyset)] \longrightarrow E'[\mathbb{H}(\Gamma', b, \emptyset)] = \mathbb{H}(\Gamma, E[b], \emptyset)$$

We then proceed by case analysis on the head redex of $a \circ \rightarrow b$.

If $\pi_i(a_1, a_2) \circ \rightarrow a_i$ this is immediate by definition of $\mathbb{H}(\Gamma', _, S)$.

If $(\lambda(x) a) b \circ \rightarrow a[b/x]_\emptyset$ then we need to prove that $\mathbb{H}(\Gamma, a[b/x]_\emptyset, \emptyset) = \mathbb{H}((\Gamma, x : \emptyset), a, \emptyset)[\mathbb{H}(\Gamma, b, \emptyset)/x]_\emptyset$. This is given by proving, by immediate induction on a , the following strengthened statement

$$\forall S, \quad \mathbb{H}(\Gamma, a[b/x]_S, S) = \mathbb{H}((\Gamma, x : \emptyset), a, S)[\mathbb{H}(\Gamma, b, \emptyset)/x]_S$$

Finally, if $\delta(\diamond, \phi.b) \circ \rightarrow b[\diamond/\phi]$, we have $\mathbb{H}(\Gamma', \delta(\diamond, \phi.b), S) = \delta(\diamond, \phi.\mathbb{H}(\Gamma', b, S \cup \{\phi\})) \circ \rightarrow \mathbb{H}(\Gamma', b, S \cup \{\phi\})[\diamond/\phi]$ and it suffices to check by direct induction on b that

$$\mathbb{H}(\Gamma', b[\diamond/\phi], S) = \mathbb{H}(\Gamma', b, S \cup \{\phi\})[\diamond/\phi]$$

□

Soundness Given a well-typed F_{th} term a , its hide-normalized form $\mathbb{H}(a)$ is still well-typed and has the same reduction behavior – errors included. We can compute the maximal hide-extrusion a' of $\mathbb{H}(a)$; this term is well-typed in both F_t and F_{th} . All that remains, to establish that the original term a is sound, is to forward-simulate any reduction path starting from a' in F_t . This should be done carefully, however, as it is *not* the case that the hide-extrusion of $\mathbb{H}(a)$ is itself hide-normal; it is, except on the subterms created by hide-extrusion: these introduced linearly-used variables to preserve scoping, and of course did not insert the appropriate hiding constructs, as its goal is to remove hiding. Then hide-extrusion may not terminate, as a result of iterating hide-normalization. Fortunately, the non-hide-normalized parts produced by hide-extrusion remain well-separated from other subterms during reduction, and are not affected by β -reduction from other parts of the term.

Theorem 3 (Soundness of F_{th}). *Every well-typed F_{th} term is sound.*

Proof. We will refer in this proof to the definitions and properties stated in the proof of reduction of hide-normalized terms 14.

Multi-extrusion We will first reformulate hide-extrusion using multi-hole contexts instead of single-hole contexts. If $E[\square_i]^i$ is a multi-hole context, let us define $E[i]$ as $E[\square/\square_i]$: this can be seen as a single-context context in the variable \square , which is in the i -th position of $E[_]^i$ (and the other \square_j for $j \neq i$ at their usual position). In particular, $\mathbf{abs}(E[i], _)$ and $\mathbf{app}(_, E[i])$ are well-defined.

We can then give the following re-definition of hide-extrusion, using $\mathbf{let} (x_i)^i = (a_i)^i$ in b as a multi-let syntax with arbitrary ordering.

EXTRUDE

$$\frac{\phi \notin E[\square_i]^{i \in I}}{\delta(b, \phi.E[\mathbf{hide} \phi \mathbf{in} a_i]^i) \rightsquigarrow \mathbf{let} (x_i)^i = (\mathbf{abs}(E[i], a_i))^i \mathbf{in} \delta(b, \phi.E[\mathbf{app}(x_i, E[i])]^i)}$$

If we index our holes with $I = \{1, 2, \dots, n\}$, immediate induction on $i \in I$ shows that the relation (\rightsquigarrow) is included in the n -th iterated hide-extrusion (\rightsquigarrow^n) . In particular, (\rightsquigarrow) and (\rightsquigarrow^n) have the same normal forms.

We will note $\mathbf{X}(a)$ the normal form of applying (\rightsquigarrow) to a – including in depth, under any term context.

Forward Simulation We now show that if $a \longrightarrow b$ among hide-normalized \mathbf{F}_{th} terms, then $\mathbf{X}(a) \longrightarrow^* \mathbf{X}(b)$ in \mathbf{F}_t , modulo some extra linear reductions.

If we have a reduction

$$\frac{a \circ \rightarrow b \quad \mathbf{unguarded}(E)}{E[a] \longrightarrow E[b]}$$

notice that a may not contain a variable ϕ bound in E , as we have $\mathbf{unguarded}(E)$; this means that E and its hole can be extruded independently: for any a' we have $\mathbf{X}(E[a']) = \mathbf{X}(E[\square])[\mathbf{X}(a')/\square]$. It then suffices to show that if $a \circ \rightarrow b$ then $\mathbf{X}(a) \longrightarrow^* \mathbf{X}(b)$ – notice that we obtain a repeated non-head reduction, as reducing extruded form may require extra work not in head position.

The $\pi_i(a_1, a_2) \circ \rightarrow a_i$ case is immediate: $\mathbf{X}(\pi_i(a_1, a_2)) = \pi_i(\mathbf{X}(a_1), \mathbf{X}(a_2)) \circ \rightarrow \mathbf{X}(a_i)$ and $\Gamma \Vdash \pi_i(a_1, a_2) : S$ implies $\Gamma \Vdash \pi_i(\mathbf{X}(a_1), \mathbf{X}(a_2)) : S$

In the case $(\lambda(x) a) b \circ \rightarrow a[b/x]_\emptyset$, we use the fact that our input term was hide-normalized: the reduction is in fact a pure substitution $(\lambda(x) a) b \circ \rightarrow a[b/x]$. In \mathbf{F}_t we have $\mathbf{H}((\lambda(x) a) b) = (\lambda(x) \mathbf{H}(a)) \mathbf{H}(b) \circ \rightarrow \mathbf{H}(a)[\mathbf{H}(b)/x] = \mathbf{H}(a[b/x])$.

The last \mathbf{F}_{th} redex to consider is $\delta(\diamond, \phi.b) \longrightarrow b[\diamond/\phi]$. We will write it in expanded form: b is of the form $E[\mathbf{hide} \phi \mathbf{in} a_i]^i$ with $\phi \notin E$, so we have $\delta(\diamond, \phi.E[\mathbf{hide} \phi \mathbf{in} a_i]^i) \longrightarrow E[a_i]^i$. The source of this reduction (\rightsquigarrow) -normalizes to $\mathbf{let} (x_i)^i = (\mathbf{abs}(E[i], \mathbf{H}(a_i)))^i \mathbf{in} \delta(\diamond, \phi.\mathbf{H}(E)[\mathbf{app}(x_i, E[i])]^i)$. We conclude by the following reduction sequence:

$$\begin{aligned} & \mathbf{let} (x_i)^i = (\mathbf{abs}(E[i], \mathbf{H}(a_i)))^i \mathbf{in} \delta(\diamond, \phi.\mathbf{H}(E)[\mathbf{app}(x_i, E[i])]^i) \\ \circ \rightarrow & \mathbf{let} (x_i)^i = (\mathbf{abs}(E[i], \mathbf{H}(a_i)))^i \mathbf{in} \mathbf{H}(E)[\mathbf{app}(x_i, E[i])]^i \\ \rightarrow^* & \mathbf{H}(E)[\mathbf{app}(\mathbf{abs}(E[i], \mathbf{H}(a_i)), E[i])]^i \\ \rightarrow^* & \mathbf{H}(E)[\mathbf{H}(a_i)]^i \\ \simeq_{\circ} & \mathbf{H}(E[a_i])^i \end{aligned}$$

Soundness If a is well-typed in F_{th} , then $H(a)$ is well-typed as well (Lemma 12), and thus $X(H(a))$ is well-typed in F_t (Lemma 9). But then, if a had a reduction path to an error in F_{th} , $H(a)$ would have one as well, by forward simulation (Lemmas 15) and preservation of errors (Lemma 13), whose intermediary reducts would be in normalized form. As $X(_)$ is a forward simulation on hide-normalized forms (we just proved this) and preserves errors (Lemma 10), this means that $X(H(a))$ would have a reduction path to an error in F_t , which is absurd by soundness of F_t (Theorem 4). □

4 Related and Future Work

Related Work

Confluence and weak reduction It appears to be folklore that there are three ways to get confluence in a weak reduction setting. One solution is to also reduce under weak binders of subterms that do not use the bound variables [Çağman and Hindley, 1998]; we cannot apply this method in F_{th} as *uses* of propositions are not traced in terms. Another solution is to introduce explicit weakening when substituting under a binder, so as to preserve the non-dependency with bound variables. This corresponds to our hiding substitution. Finally, one may use explicit substitutions and forbid them from going through weak bindings, so that the substituted terms remain reducible. Interestingly, this happens to be precisely the computational behavior of terms used in our final soundness proof (from F_{th} to F_t), as a result of hide-normalization followed by hide-extrusion.

Some explicit substitution calculi [Kesner, 2007] also have explicit weakening for the purpose of understanding reduction behavior of substructural systems (*e.g.* linear proof nets) where weakening must be applied *maximally* and this invariant is preserved by reductions and substitutions. This gives a reduction semantics that is different from our more relaxed system.

Another system with explicit weakening is **Adbmal** by Hendriks and van Oostrom [2003]. Their weakening construct enforces a well-parenthesized introduction-weakening order by removing not just one variable from scope, but also all variables introduced after it. Our hiding allows non-bracketed introduction-hiding sequences, which is more convenient for the programmer. Interestingly, we also considered a `hide * in_` construct that would hide *all* propositional variables in scope (to define hide-introducing substitution), but our use of hide-normalization, locally in the soundness proof, suffices to get the desired benefits of this batch-hiding construction. Also related to our hide-normalization technique is the *scope-extrusion* performed before **Adbmal**'s β -reductions, which extrudes the weakening above a bound variable to also weaken its binder.

System FC The family of works on System FC Sulzmann et al. [2007] is related to consistent coercion calculi in general, but also to our specific focus on implicit *v.s.* explicit use of potentially-inconsistent propositions. System FC designers

argued that explicit coercions often simplified understanding of compiler transformations, by turning semantically incorrect hard-to-debug optimizations into scope-breaking transformations that are immediately detected— see *e.g.* [Sulzmann et al., 2007, §3.8]. Implicit use is justified by user convenience, and not necessary for an intermediate compiler language; yet we claim that F_{th} provides some of these advantages in an implicit setting. The explicit reduction-blocking elimination reifies the semantic boundary into the syntax. This simplifies reasoning, for the compiler but also for the user. Another relation to our work is the march towards richer kind systems. F_{cc} included the minimum features to demonstrate usefulness, but the features studied for System FC, which moves towards a fully dependent type and kind sublanguage [Weirich et al., 2013], would make sense in our setting. In particular, with dependent kinds it should be natural to include propositions directly as kinds, allowing for example to merge product kinds, refinement kinds, and proposition conjunction as a single dependent product constructor. Consistency is known to be a pain point in the metatheory of System FC. It is neither needed nor traced in arbitrary coercion abstractions — they are not quite erasable as coercion abstraction blocks reduction. Yet, it is required for the axioms introduced at the toplevel — *e.g.* to model type families. An F_{cc} -inspired, more explicit treatment of consistency may structure System FC and provide optimization opportunities — we know that the mode of use of coercions corresponding to bounded quantification is consistent and can be erased; but the practical question of how to decide consistency is not answered in our work.

Future Work

Completing consistent coercion calculi In the process of our work we have encountered small glitches in F_{cc} : rules that we would expect to be derivable, and that were not in the current system. We have fixed them as necessary for F_{th} 's need, but some aspects could still be improved — adding η -expansions in the kind equality, and understanding whether the context consistency requirement of coercions could be removed, and recovered by a semantics argument.

Extraction Coq's extraction process Letouzey [2004] compiles a language with full reduction and explicit uses of hypotheses into OCaml, a language with weak reduction and implicit uses of hypotheses; F_{th} might be a good intermediate language in which to express and study some of the optimizations happening during the translation—which is known to be difficult. More generally, the dependent type community is aware [Brady et al., 2003] that computation is very different under arbitrary contexts. We suspect that context consistency could be a good generalization of the “empty context” assumption. A distinction between *propositional* and *definitional* truths naturally arises in our framework and, interestingly, we have a use for abstracting over definitional truths — while dependent systems don't generally consider abstracting over definitional equalities.

Conclusion

We have introduced F_{th} , a consistent coercion calculus that blocks reductions under implicit inconsistent assumptions in a fine-grained manner. This solves both practical issues (user control over reducibility) and theoretical issues (confluence) with a previous calculus of erasable coercions, F_{cc} , and opens interesting perspectives on the study of full-reduction calculi for programming language design, the interplay between type systems and weak reduction strategies, and an explicit handling of consistency in dependent type systems.

Bibliography

- E. Brady, C. McBride, and J. McKinna. Inductive families need not store their indices. In *TYPES*, pages 115–129, 2003.
- N. Çağman and J. R. Hindley. Combinatory weak reduction in lambda calculus. *Theoretical Computer Science*, 198(1):239–247, 1998.
- L. Cardelli. An implementation of FSub. Research Report 97, , 1993.
- J. Cretin. *Erasable coercions: a unified approach to type systems*. PhD thesis, Université Paris-Diderot, Paris 7, 2014.
- J. Cretin and D. Rémy. System F with Coercion Constraints. In *Logics In Computer Science (LICS)*. ACM, July 2014.
- B. Grégoire and X. Leroy. A compiled implementation of strong reduction. *ACM SIGPLAN Notices*, 37(9):235–246, 2002.
- D. Hendriks and V. van Oostrom. adbmal. In *CADE*, 2003.
- D. Kesner. The theory of calculi with explicit substitutions revisited. In *Computer Science Logic*, pages 238–252. Springer, 2007.
- D. Le Botlan and D. Rémy. MLF: Raising ML to the power of System-F. In *ICFP’80*, Aug. 2003.
- P. Letouzey. A New Extraction for Coq. In *TYPES 2002*, Feb. 2004.
- J. C. Mitchell. Polymorphic type inference and containment. *Information and Computation*, 2/3(76), 1988.
- B. Nordström, K. Petersson, and J. M. Smith. *Programming in Martin-Löfs type theory*, volume 200. Oxford University Press Oxford, 1990.
- R. Pollack. Polishing up the tait-martin-lf proof of the church-rosser theorem, 1995.
- F. Pottier. A versatile constraint-based type inference system. *Nordic Journal of Computing*, 7(4):312–347, Nov. 2000.
- V. Simonet and F. Pottier. A constraint-based approach to guarded algebraic data types. *TOPLAS*, 29(1), Jan. 2007.
- M. Sulzmann, M. M. T. Chakravarty, S. L. P. Jones, and K. Donnelly. System f with type equality coercions. In *TLDI*, pages 53–66, 2007.
- M. Takahashi. Parallel reductions in lambda-calculus. *Inf. Comput.*, 118(1): 120–127, 1995.
- D. Vytiniotis and S. P. Jones. Practical aspects of evidence-based compilation in system FC. , 2011.
- S. Weirich, J. Hsu, and R. A. Eisenberg. System FC with explicit kind equality. In *ICFP*, pages 275–286, 2013.
- H. Xi. Dependent ml an approach to practical programming with dependent types. *J. Funct. Program.*, 17(2):215–286, 2007.

$$\begin{array}{c}
\star \rightsquigarrow \{- : \star \mid \top\} \qquad 1 \rightsquigarrow \{- : 1 \mid \top\} \\
\\
\frac{\kappa_1 \rightsquigarrow \{\alpha_1 : \kappa'_1 \mid P_1\} \quad \kappa_2 \rightsquigarrow \{\alpha_2 : \kappa'_2 \mid P_2\}}{\kappa_1 * \kappa_2 \rightsquigarrow \{\alpha : \kappa'_1 * \kappa'_2 \mid P_1[\pi_1 \alpha / \alpha_1] \wedge P_2[\pi_2 \alpha / \alpha_2]\}} \\
\\
\frac{\kappa \rightsquigarrow \{\beta : \kappa' \mid Q\}}{\{\alpha : \kappa \mid P\} \rightsquigarrow \{\alpha : \kappa' \mid P \wedge Q[\alpha / \beta]\}}
\end{array}$$

Fig. 15. Splitting judgment ($\kappa \rightsquigarrow \{\alpha : \kappa' \mid P\}$)

A Additional Material

A.1 Translating F_{cc} into F_t

The advantage of propositional truths over F_{cc} 's inconsistent abstraction is that while F_{cc} 's ∂a blocks the whole term which is under the inconsistent assumption, F_{th} 's construction allows to first abstract over the assumption, and only later block subterms. It is immediate that this is more general: one can write $\lambda(x) \delta(x, \phi.a)$ to abstract over an assumption and immediately block on it.

This intuition gives a translation from F_{cc} into F_t . While it is not needed to establish soundness of F_t , it is an interesting proof because it highlight important design differences between both languages.

Splitting kinds The translation outlined above only works when F_{cc} 's inconsistent abstraction is used to abstract over a kind of the form $\langle P \rangle$ (§2.2) rather than any possibly-uninhabited kind κ . In the general case, we need to split κ into an always-consistent kind κ' , representing the computational structure of types inhabiting κ , and an inconsistent proposition P , encapsulating the logical restrictions on κ 's inhabitants. The inconsistent abstraction type $\Pi(\alpha : \kappa) \tau$ can then be translated into $\forall(\alpha : \kappa') [P] \rightarrow^b \sigma$, using consistent abstraction for the consistent part of κ and an administrative λ^b -abstraction on a proof witness of the inconsistent part.

For exactly the same reason as for the previous translation we target an administrative variant of F_t , called F_t^b , that duplicates the function type into an incompatible variant.

Given that refinement types are special cases of (dependent) products, we may see the splitting of a kind κ in a consistent kind κ' and a logical part P as the standardization of κ into a refinement type $\{\alpha : \kappa' \mid P\}$; this transformation is correct if it preserves the set of types inhabiting the kind.

We define in Figure 15 the splitting relation $\kappa \rightsquigarrow \{\alpha : \kappa' \mid P\}$. It is given as an inductive relation, but direct induction shows that κ', P are (total) functions of κ . Although F_{cc} 's kind language only has various forms of products, splitting would also work with a richer kind language. For example, if we had arrow kinds (the kind of F_ω type-level functions), we would extend the definition with this

additional rule:

$$\frac{\kappa_1 \rightsquigarrow \{\alpha_1 : \kappa'_1 \mid P_1\} \quad \kappa_2 \rightsquigarrow \{\alpha_2 : \kappa'_2 \mid P_2\}}{(\kappa_1 \rightarrow \kappa_2) \rightsquigarrow \{\alpha : \kappa'_1 \rightarrow \kappa'_2 \mid \forall(\beta : \{\beta : \kappa'_1 \mid P_1[\beta/\alpha_1]\}) P_2[(\alpha \beta)/\alpha_2]\}}$$

which shows that it extends to contravariant kind constructors.

Definition 1 (Kind equivalence). We define kind equivalence $\kappa \equiv \kappa'$ as $\forall \Gamma, \forall \tau, (\Gamma \vdash \tau : \kappa \iff \Gamma \vdash \tau : \kappa')$.

We may then show that the splitting behaves as expected: the returned kind is always consistent, and the refinement kind as a whole has the same inhabitants as the original kind.

Lemma 16 (Correctness of splitting). If $\kappa \rightsquigarrow \{\alpha : \kappa' \mid \phi\}$ holds, then $\emptyset \vdash \sigma : \kappa'$ for some σ and $\kappa \equiv \{\alpha : \kappa' \mid \phi\}$.

Proof. The proof is a straightforward induction over the derivation of the decomposition. To each derivation of $\kappa \rightsquigarrow \{\alpha : \kappa' \mid \phi\}$, we associate a derivation of the form $\emptyset \vdash \sigma : \kappa'$ for some σ and, assuming that $\Gamma \vdash \tau : \kappa$, a derivation of $\Gamma \vdash \tau : \{\alpha : \kappa' \mid \phi\}$.

$$\begin{aligned}
\star \rightsquigarrow \{- : \star \mid \top\} &\Longrightarrow \left(\frac{}{\emptyset \vdash 1 : \star}, \frac{\Gamma \vdash \tau : \star \quad \overline{\Gamma \vdash \top}}{\Gamma \vdash \tau : \{- : \star \mid \top\}} \right) \\
1 \rightsquigarrow \{- : 1 \mid \top\} &\Longrightarrow \left(\frac{}{\emptyset \vdash () : 1}, \frac{\Gamma \vdash \tau : 1 \quad \overline{\Gamma \vdash \top}}{\Gamma \vdash \tau : \{- : 1 \mid \top\}} \right) \\
\frac{\kappa_1 \rightsquigarrow \{\alpha_1 : \kappa'_1 \mid P_1\} \quad \kappa_2 \rightsquigarrow \{\alpha_2 : \kappa'_2 \mid P_2\}}{\kappa_1 * \kappa_2 \rightsquigarrow \{\alpha : \kappa'_1 * \kappa'_2 \mid P_1[\pi_1 \alpha / \alpha_1] \wedge P_2[\pi_2 \alpha / \alpha_2]\}} &\Longrightarrow \\
\left(\frac{\frac{\frac{\frac{\vdash \sigma_1 : \kappa'_1 \quad \vdash \sigma_2 : \kappa'_2}{\vdash (\sigma_1, \sigma_2) : \kappa'_1 * \kappa'_2}}{\Gamma \vdash \tau : \kappa_1 * \kappa_2}}{\Gamma \vdash \pi_i \tau : \kappa_i \equiv \{\alpha_i : \kappa'_i \mid P_i\}} \quad \frac{\frac{\frac{\Gamma \vdash \tau : \kappa_1 * \kappa_2}{\Gamma \vdash \pi_i \tau : \kappa_i \equiv \{\alpha_i : \kappa'_i \mid P_i\}}}{(\Gamma \vdash \pi_i \tau : \kappa'_i)_{i \in \{1,2\}}} \quad \frac{\frac{\frac{\Gamma \vdash \tau : \kappa_1 * \kappa_2}{\Gamma \vdash \pi_i \tau : \kappa_i} \quad \overline{\Gamma \vdash \top}}{\Gamma \vdash \pi_i \tau : \kappa_i \equiv \{\alpha_i : \kappa'_i \mid P_i\}}}{(\Gamma \vdash P_i[\pi_i \tau / \alpha_i])_{i \in \{1,2\}}} \quad \frac{\Gamma \vdash \tau : \kappa_1 * \kappa_2}{\Gamma \vdash P_1[\pi_1 \tau / \alpha_1] \wedge P_2[\pi_2 \tau / \alpha_2]}}{\Gamma \vdash \tau : \{\alpha : \kappa'_1 * \kappa'_2 \mid P_1[\pi_1 \alpha / \alpha_1] \wedge P_2[\pi_2 \alpha / \alpha_2]\}} \right)
\end{aligned}$$

$$\frac{\kappa \rightsquigarrow \{\beta : \kappa' \mid Q\}}{\{\alpha : \kappa \mid P\} \rightsquigarrow \{\alpha : \kappa' \mid P \wedge Q[\alpha / \beta]\}} \Longrightarrow$$

$$\left(\frac{\frac{\frac{\frac{\vdash \sigma : \kappa'}{\Gamma \vdash \tau : \{\alpha : \kappa \mid P\}}}{\Gamma \vdash \tau : \kappa \equiv \{\beta : \kappa' \mid Q\}}}{\Gamma \vdash \tau : \kappa'} \quad \frac{\frac{\frac{\Gamma \vdash \tau : \{\alpha : \kappa \mid P\}}{\Gamma \vdash \tau : \{\alpha : \kappa \mid P\}} \quad \frac{\Gamma \vdash \tau : \{\alpha : \kappa \mid P\}}{\Gamma \vdash P[\tau / \alpha]}}{\Gamma \vdash P[\tau / \alpha]} \quad \frac{\frac{\Gamma \vdash \tau : \{\alpha : \kappa \mid P\}}{\Gamma \vdash \tau : \kappa \equiv \{\beta : \kappa' \mid Q\}}}{\Gamma \vdash Q[\tau / \beta]}}{\Gamma \vdash Q[\tau / \beta]}}{\Gamma \vdash P[\tau / \alpha] \wedge Q[\tau / \beta]} \quad \frac{\Gamma \vdash \tau : \{\alpha : \kappa \mid P\}}{\Gamma \vdash \tau : \{\alpha : \kappa \mid P\}}}{\Gamma \vdash \tau : \{\alpha : \kappa' \mid P \wedge Q[\alpha / \beta]\}} \right)$$

□

η-expansion of product kinds To prove correctness of splitting in the case $\kappa_1 * \kappa_2$, we used the admissible rule

$$\frac{\Gamma \vdash \pi_1 \tau : \kappa_1 \quad \Gamma \vdash \pi_2 \tau : \kappa_2}{\Gamma \vdash \tau : \kappa_1 * \kappa_2}$$

Unfortunately, this rule is not derivable in F_{cc} – or F_{th} . We proved its admissibility by adding it to the mechanized semantics of F_{cc} and updating the machine-checked soundness proof accordingly.⁶

This rule ought to be derivable rather than merely admissible. The best way to do this would be to add η -expansion of product kinds to the kind equality: $\tau =_{(\kappa_1 * \kappa_2)} (\pi_1 \tau, \pi_2 \tau)$. But the kind equality is currently untyped, and it should be made typed to conveniently interpret η -expansion. This is a possibly invasive change to the F_{cc} metatheory, independent of our present purpose, that we left for future work.

The translation We now define the translation $\llbracket - \rrbracket$ from F_{cc} to F_t^b , the variant of F_t with administrative arrows:

$$\begin{aligned} \llbracket \Pi(\alpha : \kappa) \tau \rrbracket &\triangleq \forall(\alpha : \llbracket \kappa' \rrbracket) \llbracket \llbracket P \rrbracket \rrbracket \rightarrow^b \llbracket \tau \rrbracket \quad (\kappa \rightsquigarrow \{\alpha : \kappa' \mid P\}) \\ \llbracket \partial a \rrbracket &\triangleq \lambda^b(x) \delta(x, \phi, \llbracket a \rrbracket) \quad (\phi \text{ fresh}) \\ \llbracket a \diamond \rrbracket &\triangleq \llbracket a \rrbracket^b \diamond \end{aligned}$$

On the well-kindedness check in F_{cc} and F_t We encounter a difficulty when proving that the translation preserves type information. Consider a F_{cc} term ∂a at type $\llbracket \Pi(\alpha : \kappa) \tau \rrbracket$, with $\kappa \rightsquigarrow \{\alpha : \kappa' \mid P\}$. Its F_t^b translation is $\lambda^b(x) \delta(x, \phi, \llbracket a \rrbracket)$. When type-checking the F_{cc} term, its subterm a is type-checked in the context $(\Gamma, \alpha : \kappa)$. In the translated term, the translated subterm $\llbracket a \rrbracket$ appears in the context $(\llbracket \Gamma \rrbracket, \alpha : \kappa', x : \llbracket \llbracket P \rrbracket \rrbracket, \phi : \llbracket \llbracket P \rrbracket \rrbracket)$. It is rather clear that this translated context contains the same static information as the original context: the pair of assumptions $(\alpha : \llbracket \kappa' \rrbracket, \phi : \llbracket \llbracket P \rrbracket \rrbracket)$ is interchangeable with $(\alpha : \llbracket \{\alpha : \kappa' \mid P\} \rrbracket)$, itself interchangeable with $(\alpha : \llbracket \kappa \rrbracket)$.

Lemma 17. *If $\kappa \rightsquigarrow \{\alpha : \kappa' \mid P\}$, the F_{cc} contexts $(\Gamma, \alpha : \kappa)$ and $(\Gamma, \alpha : \kappa', \phi : P)$ prove the same judgments.*

Proof. We have $\frac{\Gamma, \alpha : \kappa', \phi : P \vdash \alpha : \kappa' \quad \Gamma, \alpha : \kappa', \phi : P \vdash P}{\Gamma, \alpha : \kappa', \phi : P \vdash \alpha : \{\alpha : \kappa' \mid P\}}$ so, by correctness (Lemma 16) of the split $\kappa \rightsquigarrow \{\alpha : \kappa' \mid P\}$, we also have $\Gamma, \alpha : \kappa', \phi : P \vdash \alpha : \kappa$. By substitution, any judgment of the form $\Gamma, \beta : \kappa \vdash J$ can therefore be transformed into a valid judgment $\Gamma, \alpha : \kappa', \phi : P \vdash J[\alpha/\beta]$. □

Now comes the difficulty: while the context $(\llbracket \Gamma \rrbracket, \alpha : \llbracket \kappa' \rrbracket, \phi : \llbracket \llbracket P \rrbracket \rrbracket)$ is available when type-checking the term-level translation of ∂a , we get a *strictly weaker* context when checking the well-formedness of translated *types*: checking $\Gamma \vdash \Pi(\alpha : \kappa) \tau : \star$ recursively checks the subtype τ in the context $\Gamma, \alpha : \kappa$, but checking the translation $\forall(\alpha : \llbracket \kappa' \rrbracket) \llbracket \llbracket P \rrbracket \rrbracket \rightarrow^b \llbracket \tau \rrbracket$ processes the subtype τ in the weaker context $\llbracket \Gamma \rrbracket, \alpha : \llbracket \kappa' \rrbracket, x : \llbracket \llbracket P \rrbracket \rrbracket$: P is not available as an implicit

⁶ The modified Coq sources are available at <https://github.com/gasche/fcc/tree/eta-expansion-on-product-kinds>.

proposition, only as a dynamic witness – in particular, it is not usable as an hypothesis for proposition satisfiability judgments.

This is not an innocuous detail, but one that comes from an important design difference between F_{cc} and F_t . In F_{cc} , the inconsistent abstraction always block the reduction of the whole term it types, so the type system can statically assume that the (possibly inconsistent) abstracted kind is implicitly available when type-checking the term – or checking well-formedness of its type, etc.

On the contrary, in F_t blocking may happen only in subterms or not at all, so no such static assumption is possible. In the translation we use, any λ on a propositional truth $[P]$ is immediately followed by an elimination form that adds the proposition P in the context – matching the F_{cc} typing; but this term-level invariant is not visible in the type itself. So the type-checking judgment (for terms) has the same assumptions than the source F_{cc} derivation, but the well-formedness (kinding) judgment for types has to work under weaker assumptions.

Fortunately, this weaker environment is nonetheless enough to check type well-formedness. The idea is that while type-checking terms is about *semantic* aspects of correctness that may require to check that propositions are satisfiable and contexts inhabited, well-formedness on the other hand is a syntactic criterion that does not depend on the ability of the context to prove propositions. This is justified by the following non-trivial lemma.

Lemma 18 (Phase separation in F_{cc}). *The F_{cc} contexts $\Gamma, \phi : P$ and $\Gamma, \phi : \top$ prove the same star-kinding ($\Gamma \vdash \tau : \star$) and well-formedness judgments ($\Gamma \Vdash t$) – for t ranging over all environments, kinds, propositions and coercions.*

Notice that the judgments differ in general for type-checking terms ($\Gamma \vdash a : \tau$), kinding inconsistent kinds ($\Gamma \vdash \tau : \kappa$), proving coercions ($\Gamma \vdash (\Sigma \vdash \tau) \triangleright \sigma$), satisfiability of propositions ($\Gamma \vdash P$) and consistency of contexts ($\Gamma \vdash \exists \Delta$).

Proof. The proof, by mutual induction, is immediate for the well-formedness judgments: all premises of a well-formedness rule are well-formedness judgments themselves, expect some star-kinding hypothesis $\Gamma \vdash \tau : \star$. It is not immediate, however, for the star-kinding judgment itself, due to kinding rules such that may require proving arbitrary propositions, such as

$$\frac{\Gamma \vdash \tau : \kappa \quad \Gamma \vdash P[\tau/\alpha]}{\Gamma \vdash \tau : \{\alpha : \kappa \mid P\}}$$

The intuition for finishing the proof is to notice that such a rule does not classifies types of kind \star , and that in fact it should never be needed inside the derivation of a star-kinding judgment. To formally establish this, we need to introduce a notion of *normal form* for kinding proofs – which would correspond to a beta-normal form on their explicit term representation. Let us define that the inference rules for variable and unit are *base rules*, the inference rules for projection out of a product kind or out of a refinement kind are *destructor rules*, and that other rules, expect the variable rules, are *constructor rules* – including the conversion rule. A proof in *normal form* has base rules followed by zero or

one equality rule, followed by a series of destructor rules, followed by a series of constructor rules mixed with conversion rules – constructor and destructor rules do not alternate.

We claim that any derivation for a judgment $\Gamma \vdash \tau : \kappa$ can be rewritten into a derivation in normal form for the same judgment. Compressing several consecutive conversion rules in at most one is immediate, by transitivity of equality. We will first show that any equality after a destructor can be moved before it. We will then show that any pair of a destructor rule below a constructor rule (with optionally a conversion between them) can be removed. Repeating those transformations gives a derivation in normal form.

Conversion after destructor

$$\frac{\frac{\Gamma \vdash \tau : \kappa_1 * \kappa_2}{\Gamma \vdash \pi_1 \tau : \kappa_1} \quad \kappa_1 =_\beta \kappa'_1}{\Gamma \vdash \pi_1 \tau : \kappa'_1} \Longrightarrow \frac{\frac{\Gamma \vdash \tau : \kappa_1 * \kappa_2 \quad \kappa_1 * \kappa_2 =_\beta \kappa'_1 * \kappa_2}{\Gamma \vdash \tau : \kappa'_1 * \kappa_2}}{\Gamma \vdash \pi_1 \tau : \kappa'_1}$$

$$\frac{\Gamma \vdash \tau : \{\alpha : \kappa \mid P\}}{\Gamma \vdash \tau : \kappa} \quad \kappa =_\beta \kappa'}{\Gamma \vdash \tau : \kappa'} \Longrightarrow \frac{\frac{\Gamma \vdash \tau : \{\alpha : \kappa \mid P\} \quad \{\alpha : \kappa \mid P\} =_\beta \{\alpha : \kappa' \mid P\}}{\Gamma \vdash \tau : \{\alpha : \kappa' \mid P\}}}{\Gamma \vdash \tau : \kappa'}$$

Constructor-destructor pair If a destructor rule for the product kind is immediately after a constructor rule, we know that the constructor can only be the product kind constructor – types have to match. If a conversion rule is placed between them, it must be for an equality of the form $\kappa'_1 * \kappa'_2 =_\beta \kappa_1 * \kappa_2$, respecting the product structure. Those can be decomposed into a pair of equalities $\kappa'_1 =_\beta \kappa_1$ and $\kappa'_2 =_\beta \kappa_2$. Note there are no equalities between kinds with different head constructors (eg. product and refinement), so in-between equalities do not break the constructor-destructor correspondence.

The case of refinement kinds is similar: if there is an equality, it is of the form $\{\alpha : \kappa' \mid P'\} = \{\alpha : \kappa \mid P\}$, which implies both equalities $\kappa' = \kappa$ and $P' = P$.

We will give the rewrites for the case where a conversion is present. Just erasing the conversions gives the valid rewrite without.

$$\frac{\frac{\Gamma \vdash \tau_1 : \kappa'_1 \quad \Gamma \vdash \tau_2 : \kappa'_2}{\Gamma \vdash (\tau_1, \tau_2) : \kappa'_1 * \kappa'_2} \quad \kappa'_1 * \kappa'_2 =_\beta \kappa_1 * \kappa_2}{\frac{\Gamma \vdash (\tau_1, \tau_2) : \kappa_1 * \kappa_2}{\Gamma \vdash \pi_i(\tau_1, \tau_2) : \kappa_i}} \Longrightarrow \frac{\Gamma \vdash \tau_i : \kappa'_i \quad \kappa'_i = \kappa_i}{\Gamma \vdash \tau_i : \kappa_i}$$

$$\frac{\frac{\Gamma \vdash \tau : \kappa' \Gamma \vdash P'[\tau/\alpha]}{\Gamma \vdash \{\alpha : \kappa' \mid P'\}} \quad \{\alpha : \kappa' \mid P'\} =_\beta \{\alpha : \kappa \mid P\}}{\frac{\Gamma \vdash \tau : \{\alpha : \kappa \mid P\}}{\Gamma \vdash \tau : \kappa}} \Longrightarrow \frac{\Gamma \vdash \tau : \kappa' \quad \kappa =_\beta \kappa'}{\Gamma \vdash \tau : \kappa}$$

Proving phase separation Now that we can rewrite derivation in normal form, we can prove phase separation for star-kinded judgments. If we have a proof of $\Gamma \vdash \tau : \star$, we can without loss of generality consider a *normal* proof of this judgment.

Such a proof cannot contain a constructor rule for a kind other than \star : it cannot end with such a rule, and an equality cannot rewrite \star into another kind constructor, so it cannot be above the last conversion rule either. Our normal proof therefore only contains conversion and destructor rules, and constructor rules at the kind \star . It is immediate to verify that those rules transitively require only star-kinding or well-formedness judgments, which conclude our proofs. \square

In fact, this phase separation property was false in the F_{cc} calculus as presented in previous work, because of a known glitch in the well-formedness condition for coercions (seen as propositions). This rule, along with the rule injecting the propositional judgment into the coercion judgment, read as follows:

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

Whereas they are in the present document in the following form:

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

The consistency of the coercion context Σ was previously enforced in the well-formedness judgment for coercions, violating the phase separation between well-formedness checking and type-checking. In the new presentation, it is not required there, but only when injecting a proposition (which is a coercion) in the coercion judgment – for the other rules of the coercion judgment, one can check that the context is consistent by construction. Of course, we have checked that this change does not affect the soundness of F_{cc} , by updating Cretin's Coq development to mechanically verify the soundness of the new presentation.⁷

We now have all the pieces to prove that the translation preserves typing.

⁷ The modified Coq sources are available at <https://github.com/gasche/fcc/tree/weaken-coercion-requirement-2>.

Lemma 19 (Typing preservation of F_{cc}). *If $\Gamma \vdash a : \tau$ in F_{cc} , then $\llbracket \Gamma \rrbracket \vdash \llbracket a \rrbracket : \llbracket \tau \rrbracket$ in F_t^b .*

Proof. The translation $\llbracket _ \rrbracket$ is extended to derivations in Figure 16. Direct induction shows that the translation of a valid derivation for the F_t judgment $\Gamma \vdash a : \tau$ is a valid derivation for the F_{cc}^b judgment $\llbracket \Gamma \rrbracket \vdash \llbracket a \rrbracket : \llbracket \tau \rrbracket$ – and mutually for all judgment forms.

For convenience and readability, the derivations in this proof use the following derivable coercion rule:

$$\frac{\text{GEN-INST ETA-RULE} \quad \Gamma, \alpha_1 : \kappa_1, \Sigma \vdash \sigma : \kappa_2 \quad \Gamma, \alpha_1 : \kappa_1 \vdash (\Sigma \vdash \tau_2[\sigma/\alpha_2]) \triangleright \tau_1}{\Gamma \vdash (\Sigma \vdash \forall(\alpha_2 : \kappa_2) \tau_2) \triangleright \forall(\alpha_1 : \kappa_1) \tau_1}$$

When translating a proof that the type $\Pi(\alpha : \kappa) \tau$ is well-formed under Γ , with the splitting $\kappa \rightsquigarrow \{\beta : \kappa' \mid P\}$, one translates an hypothesis of the form $\Gamma, \alpha : \kappa \vdash \tau : \star$ into the hypothesis $\llbracket \Gamma \rrbracket, \beta : \llbracket \kappa' \rrbracket, y : \llbracket [P] \rrbracket \vdash \llbracket \tau[\beta/\alpha] \rrbracket$. This is where we need to use Lemma 18. \square

Preservation of the dynamic semantics Independently of typing preservation, we now prove the preservation of the dynamic semantics by exhibiting a bisimulation. This proof closely follow the structure we used for the translation of F_t to F_{cc}^b in §3.2, and is less detailed.

Lemma 20 (Error preservation of F_{cc}). *A term a is an error in F_{cc} if and only if $\llbracket a \rrbracket$ is an error in F_t^b .*

Lemma 21 (Bisimulation of F_t^b by F_{cc}). *Let a be a term in F_{cc} . If $a \longrightarrow a'$ for some term a' of F_{cc} , then $\llbracket a \rrbracket \xrightarrow{?}_b \cdot \longrightarrow \llbracket a' \rrbracket$. Conversely, if $\llbracket a \rrbracket \xrightarrow{?}_b \cdot \longrightarrow b'$ for some term b' of F_t^b , then there exists a term a' of F_{cc} such that $a \longrightarrow a'$ and $\llbracket a' \rrbracket = b$.*

Proof. The proof goes as the bisimulation proof for F_t 3, so we will only show the behavior of the F_{cc} -specific redex:

$$\begin{aligned} & \llbracket (\partial a) \diamond \rrbracket \\ &= (\lambda^b(x) \delta(x, \phi. \llbracket a \rrbracket))^b \diamond \\ &\circ \rightarrow_b \delta(\diamond, \phi. \llbracket a \rrbracket) \\ &\circ \rightarrow \llbracket a \rrbracket \end{aligned}$$

Conversely, only the translation of a ∂ -redex gives a λ^b -redex and then a \diamond -redex, and both reductions must happen consecutively to remain in the image of the translation. \square

Corollary 4. *If $\llbracket a \rrbracket$ is sound in F_t^b , then a is sound in F_{cc} .*

Corollary 5. *If F_t^b is sound then F_{cc} is sound.*

The justification for the relative soundness result of F_t was to establish soundness of this new calculus, from the existing proof of soundness of F_{cc} . One may wonder why we also prove that the soundness of F_t implies the soundness of F_{cc} . From a design point of view, we consider that F_{th} is a better language than F_{cc} , and should be the target of future uses and extensions. In this context of an evolving F_{th} language and soundness proof, we may want to rebuild a soundness proof for F_{cc} on top of the future F_{th} proofs.

A.2 Soundness of the administrative arrow

We now conclude our soundness proof for F_t by proving that the soundness of F_{cc} implies the soundness of F_{cc}^b . We will use a general enough proof that also shows that F_t^b is sound. Let X be either F_t or F_{cc} , and X^b the corresponding administrative variant. We prove that if X is sound, then X^b is sound.

As we previously noted, a translation that simply erases the difference between the arrow and its administrative variant is not satisfying, as it does not preserve error: $(\lambda^b(x) a) b$ and $(\lambda(x) a)^b b$ are errors in X^b , and would not be anymore with the naive erasure translation.

We could decide, of course, to use a different translation for which the counter-examples above really are errors by *boxing* the underlying λ -abstraction under an incompatible constructor. For example, translating $\lambda^b(x) a$ into the pair $((, \lambda(x) a)$, and translating the application $a^b b$ with an extra projection $(\pi_2 a) b$ would be a valid translation, in which $(\lambda(x) a)^b b$ translates to $(\pi_2 (\lambda(x) a)) b$, which is an error. But the problem is only shifted around: $\pi_1 (\lambda^b(x) a)$, which was correctly translated into an error by the previous translation, now translates into a valid term. We parametrize our translation from X^b to X over any such possible boxing.

Definition 2. A boxing B is a triple $(\mathbf{box}(\square), \mathbf{pack}(\square), \mathbf{unpack}(\square))$ of a type context and two reduction contexts of F_{cc} , such that the following properties holds:

$$\begin{aligned} & \text{if } \Gamma \vdash \tau : \star \text{ then } \Gamma \vdash \mathbf{box}(\tau) : \star \\ & x : \tau \vdash \mathbf{pack}(x) : \mathbf{box}(\tau) \\ & x : \mathbf{box}(\tau) \vdash \mathbf{unpack}(x) : \tau \\ & \text{if } \Gamma \vdash \tau \triangleright \sigma \text{ then } \Gamma \vdash \mathbf{box}(\tau) \triangleright \mathbf{box}(\sigma) \\ & \mathbf{unpack}(\mathbf{pack}(a)) \longrightarrow^* a \end{aligned}$$

Besides the identity boxing $B_{\text{id}} \triangleq (\square, \square, \square)$, we have presented above a pairing boxing $B_{\pi_1} \triangleq ((1 * \square), ((, \square), (\pi_2 \square))$; we could also use a thunking boxing $B_\lambda \triangleq ((1 \rightarrow \square), (\lambda(-) \square), (\square ()))$. The soundness proof relies on the fact that there exists at least two boxings with different head constructors.

Given any boxing B , we define the translation $\llbracket _ \rrbracket_B$ from X^b to X as follows:

$$\begin{aligned} \llbracket \tau \rightarrow^b b \rrbracket_B & \triangleq \mathbf{box}(\llbracket \tau \rrbracket_B \rightarrow \llbracket b \rrbracket_B) \\ \llbracket \lambda^b(x) a \rrbracket_B & \triangleq \mathbf{pack}(\lambda(x) \llbracket a \rrbracket_B) \\ \llbracket a^b b \rrbracket_B & \triangleq \mathbf{unpack}(\llbracket a \rrbracket_B) \llbracket b \rrbracket_B \end{aligned}$$

Lemma 22 (Forward simulation of X^b in X). *If $a \longrightarrow b$ in X^b then, for any boxing B , $\llbracket a \rrbracket_B \longrightarrow^+ \llbracket b \rrbracket_B$.*

Proof. A common property of both F_{cc} and F_t is that redexes for distinct constructors do not overlap – in term of rewrite system, we use an *orthogonality* property of our languages. We therefore know that the translation leaves unchanged the non-administrative redexes. In the administrative case, we have:

$$\begin{aligned}
& \llbracket (\lambda^b(x) a)^b b \rrbracket_B \\
= & \text{unpack}(\text{pack}(\lambda(x) \llbracket a \rrbracket_B)) \llbracket b \rrbracket_B \\
\longrightarrow^* & \lambda(x) \llbracket a \rrbracket_B \llbracket b \rrbracket_B \\
\circlearrowright & \llbracket a \rrbracket_B \llbracket \llbracket b \rrbracket_B / x \rrbracket \\
= & \llbracket a[b/x] \rrbracket_B
\end{aligned}$$

□

Lemma 23 (Preservation of X^b errors in X). *If a is an error in X^b then, for some boxing B , $\llbracket a \rrbracket_B$ is an error in X .*

Proof. Errors that do not involve the administrative constructions are translated, unchanged, by any boxing B .

Consider now an error of the form $a^b b$, where the subterm a starts with a head constructor other than λ^b ; if it is a λ (a is of the form $\lambda(x) a'$), then using the pair boxing B_{π_1} we have $\llbracket a^b b \rrbracket_{B_{\pi_1}} = \pi_2(\lambda(x) \llbracket a' \rrbracket_{B_{\pi_1}}) \llbracket b \rrbracket_{B_{\pi_1}}$, which is an error as expected. If the head constructor is not a λ , then using the identity boxing (B_{id}) or lambda boxing (B_λ) translates to an error as expected; for example, $\llbracket (a_1, a_2)^b b \rrbracket_{B_\lambda} = (\llbracket a_1 \rrbracket_{B_\lambda}, \llbracket a_2 \rrbracket_{B_\lambda}) () \llbracket b \rrbracket_{B_\lambda}$ is an error as expected.

The situation is symmetric for an error of the form $D[\lambda^b(x) a]$, where an administrative lambda is placed under an elimination form that is not the administrative application: if it is an application, one should use B_{π_1} to get an error, and otherwise B_{id} or B_λ .

□

Corollary 6. *If, for all boxing B , $\llbracket a \rrbracket_B$ is sound in X , then a is sound in X^b .*

Lemma 24 (Typing preservation of X^b).

If $\Gamma \vdash a : \tau$ in X^b then, for any boxing B , we have $\llbracket \Gamma \rrbracket_B \vdash \llbracket a \rrbracket_B : \llbracket \tau \rrbracket_B$ in X .

Proof. As for the previous proofs of typing preservation, one needs to provide a derivable translation of the typing rules for the administrative arrow. The same translation applies to both F_{cc} and F_t , as it only use constructs in the intersection of the two languages.

$$\begin{aligned}
& \left[\frac{\Gamma, x : \tau \vdash a : \sigma}{\Gamma \vdash \lambda^b(x) a : \tau \rightarrow^b \sigma} \right]_B \triangleq \frac{\frac{[\Gamma]_B, x : [\tau]_B \vdash [a]_B : [\sigma]_B}{[\Gamma]_B \vdash \lambda(x) [a]_B : [\tau]_B \rightarrow [\sigma]_B}}{[\Gamma]_B \vdash \text{pack}(\lambda(x) [a]_B) : \text{box}([\tau]_B \rightarrow [\sigma]_B)}} \\
& \left[\frac{\Gamma \vdash a : \tau \rightarrow^b \sigma \quad \Gamma \vdash b : \tau}{\Gamma \vdash a^b b : \sigma} \right]_B \triangleq \\
& \frac{\frac{[\Gamma]_B \vdash [a]_B : \text{box}([\tau]_B \rightarrow [\sigma]_B)}{[\Gamma]_B \vdash \text{unpack}([a]_B) : [\tau]_B \rightarrow [\sigma]_B} \quad [\Gamma]_B \vdash [b]_B : [\tau]_B}{[\Gamma]_B \vdash \text{unpack}([a]_B) [b]_B : [\sigma]_B}} \\
& \left[\frac{\Gamma \vdash \tau : \star \quad \Gamma \vdash \sigma : \star}{\Gamma \vdash \tau \rightarrow^b \sigma : \star} \right]_B \triangleq \frac{\frac{[\Gamma]_B \vdash [\tau]_B : \star \quad [\Gamma]_B \vdash [\sigma]_B : \star}{[\Gamma]_B \vdash [\tau]_B \rightarrow [\sigma]_B : \star}}{[\Gamma]_B \vdash \text{box}([\tau]_B \rightarrow [\sigma]_B) : \star}} \\
& \left[\frac{\Gamma, \Sigma \vdash \tau' \triangleright \tau \quad \Gamma \vdash (\Sigma \vdash \sigma) \triangleright \sigma'}{\Gamma \vdash (\Sigma \vdash (\tau \rightarrow^b \sigma)) \triangleright (\tau' \rightarrow^b \sigma')} \right]_B \triangleq \\
& \frac{\frac{[\Gamma]_B, [\Sigma]_B \vdash [\tau']_B \triangleright [\tau]_B \quad [\Gamma]_B \vdash ([\Sigma]_B \vdash [\sigma]_B) \triangleright [\sigma']_B}{[\Gamma]_B \vdash ([\Sigma]_B \vdash [\tau]_B \rightarrow [\sigma]_B) \triangleright [\tau']_B \rightarrow [\sigma']_B}}{[\Gamma]_B \vdash ([\Sigma]_B \vdash \text{box}([\tau]_B \rightarrow [\sigma]_B)) \triangleright \text{box}([\tau']_B \rightarrow [\sigma']_B)}}
\end{aligned}$$

□

Corollary 7. *If X is sound, then X^b is sound.*

Theorem 4. *F_t is sound.*

Proof. Our translation of F_t into F_{cc}^b shows that if F_{cc}^b is sound, then F_t is sound (Corollary 2). We just demonstrated in turn that if F_{cc} is sound, then F_{cc}^b is sound (Corollary 7). Finally, soundness of F_{cc} is a result of previous work 1.

□

$$\begin{array}{c}
\left[\frac{\Gamma \vdash \kappa \quad \Gamma, \alpha : \kappa \vdash \tau : \star}{\Gamma \vdash \Pi(\alpha : \kappa) \tau : \star} \right] \stackrel{\kappa \rightsquigarrow \{\beta : \kappa' \mid P\}}{\triangleq} \\
\frac{\frac{\frac{\llbracket \Gamma \rrbracket, \beta : \llbracket \kappa' \rrbracket \vdash \llbracket P \rrbracket \quad \llbracket \Gamma \rrbracket, \beta : \llbracket \kappa' \rrbracket, y : \llbracket \llbracket P \rrbracket \rrbracket \vdash \llbracket \tau[\beta/\alpha] \rrbracket : \star}{\llbracket \Gamma \rrbracket, \beta : \llbracket \kappa' \rrbracket \vdash \llbracket \llbracket P \rrbracket \rrbracket \rightarrow^b \llbracket \tau[\beta/\alpha] \rrbracket : \star}}{\llbracket \Gamma \rrbracket \vdash \llbracket \kappa' \rrbracket}}{\llbracket \Gamma \rrbracket \vdash \forall(\beta : \llbracket \kappa' \rrbracket) (\llbracket \llbracket P \rrbracket \rrbracket \rightarrow^b \llbracket \tau[\beta/\alpha] \rrbracket) : \star}} \\
\left[\frac{\Gamma \vdash \kappa \quad \Gamma, \alpha : \kappa \vdash a : \tau}{\Gamma \vdash \partial a : \Pi(\alpha : \kappa) \tau} \right] \stackrel{\kappa \rightsquigarrow \{\beta : \kappa' \mid P\}}{\triangleq} \\
\frac{\frac{\frac{\llbracket \Gamma \rrbracket, \beta : \llbracket \kappa' \rrbracket, x : \llbracket \llbracket P \rrbracket \rrbracket \vdash x : \llbracket \llbracket P \rrbracket \rrbracket \quad \llbracket \Gamma \rrbracket, \beta : \llbracket \kappa' \rrbracket, x : \llbracket \llbracket P \rrbracket \rrbracket, \phi : \llbracket P \rrbracket \vdash \llbracket a \rrbracket : \llbracket \tau[\beta/\alpha] \rrbracket}{\llbracket \Gamma \rrbracket, \beta : \llbracket \kappa' \rrbracket, x : \llbracket \llbracket P \rrbracket \rrbracket \vdash \delta(x, \phi, \llbracket a \rrbracket) : \llbracket \tau[\beta/\alpha] \rrbracket}}{\llbracket \Gamma \rrbracket \vdash \lambda^b(x) \delta(x, \phi, \llbracket a \rrbracket) : \forall(\beta : \llbracket \kappa' \rrbracket) (\llbracket \llbracket P \rrbracket \rrbracket \rightarrow^b \llbracket \tau[\beta/\alpha] \rrbracket)}} \\
\left[\frac{\Gamma \vdash a : \Pi(\alpha : \kappa) \tau \quad \Gamma \vdash \sigma : \kappa}{\Gamma \vdash a \diamond : \tau[\sigma/\alpha]} \right] \stackrel{\kappa \rightsquigarrow \{\beta : \kappa' \mid P\}}{\triangleq} \\
\frac{\frac{\frac{\llbracket \Gamma \rrbracket \vdash \llbracket a \rrbracket : \forall(\beta : \llbracket \kappa' \rrbracket) (\llbracket \llbracket P \rrbracket \rrbracket \rightarrow^b \llbracket \tau[\beta/\alpha] \rrbracket) \quad \llbracket \Gamma \rrbracket \vdash \llbracket \sigma \rrbracket : \llbracket \kappa' \rrbracket}{\llbracket \Gamma \rrbracket \vdash \llbracket a \rrbracket : \llbracket \llbracket P \rrbracket \rrbracket \rightarrow^b \llbracket \tau[\beta/\alpha] \rrbracket} \llbracket \llbracket \sigma \rrbracket / \beta}}{\llbracket \Gamma \rrbracket \vdash \llbracket a \rrbracket^b \diamond : \llbracket \tau[\sigma/\alpha] \rrbracket}} \quad \frac{\llbracket \Gamma \rrbracket \vdash \llbracket P \rrbracket}{\llbracket \Gamma \rrbracket \vdash \llbracket \diamond \rrbracket : \llbracket \llbracket P \rrbracket \rrbracket}} \\
\left[\frac{\Gamma \vdash \kappa_1 \quad \Gamma \vdash \exists \Sigma \quad \Gamma, \alpha_1 : \kappa_1, \Sigma \vdash \sigma_2 : \kappa_2 \quad \Gamma, \alpha_1 : \kappa_1 \vdash (\Sigma \vdash \tau_2[\sigma_2/\alpha_2]) \triangleright \tau_1}{\Gamma \vdash (\Sigma \vdash \Pi(\alpha_2 : \kappa_2) \tau_2) \triangleright \Pi(\alpha_1 : \kappa_1) \tau_1} \right] \stackrel{(\kappa_i \rightsquigarrow \{\beta_i : \kappa'_i \mid P_i\})_i}{\triangleq} \\
\text{(GEN-INST ETA-RULE)} \\
\frac{\frac{\frac{\llbracket \Gamma \rrbracket, \beta_1 : \kappa'_1 \vdash \llbracket \llbracket P_1 \rrbracket \rrbracket \triangleright \llbracket \llbracket P_2[\sigma_2/\alpha_2] \rrbracket \rrbracket}{\llbracket \Gamma \rrbracket, \beta_1 : \kappa'_1 \vdash (\llbracket \llbracket \Sigma \rrbracket \rrbracket \vdash \llbracket \tau_2[\sigma_2/\alpha_2] \rrbracket) \triangleright \llbracket \tau_1[\beta_1/\alpha_1] \rrbracket}}{\llbracket \Gamma \rrbracket, \beta_1 : \kappa'_1, \llbracket \llbracket \Sigma \rrbracket \rrbracket \vdash \llbracket \sigma_2 \rrbracket : \llbracket \llbracket \kappa'_2 \rrbracket \rrbracket} \quad \frac{\llbracket \Gamma \rrbracket, \beta_1 : \kappa'_1 \vdash (\llbracket \llbracket \Sigma \rrbracket \rrbracket \vdash (\llbracket \llbracket P_2 \rrbracket \rrbracket \rightarrow^b \llbracket \tau_2[\beta_2/\alpha_2] \rrbracket)) \llbracket \llbracket \sigma_2 \rrbracket / \beta_2 \rrbracket) \triangleright \llbracket \llbracket P_1 \rrbracket \rrbracket \rightarrow^b \llbracket \tau_1[\beta_1/\alpha_1] \rrbracket}}{\llbracket \Gamma \rrbracket \vdash (\llbracket \llbracket \Sigma \rrbracket \rrbracket \vdash \forall(\beta_2 : \llbracket \llbracket \kappa'_2 \rrbracket \rrbracket) (\llbracket \llbracket P_2 \rrbracket \rrbracket \rightarrow^b \llbracket \tau_2[\beta_2/\alpha_2] \rrbracket)) \triangleright \forall(\beta_1 : \llbracket \llbracket \kappa'_1 \rrbracket \rrbracket) (\llbracket \llbracket P_1 \rrbracket \rrbracket \rightarrow^b \llbracket \tau_1[\beta_1/\alpha_1] \rrbracket)}}
\end{array}$$

Fig. 16. Translation of F_{cc} into F_t^b