On the Semantic Expressiveness of Recursive Types

MARCO PATRIGNANI, Stanford University, USA and CISPA Helmholtz Center for Information Security, Germany

ERIC MARK MARTIN, Stanford University, USA

DOMINIQUE DEVRIESE, Vrije Universiteit Brussel, Belgium

Recursive types extend the simply-typed lambda calculus (STLC) with the additional expressive power to enable diverging computation and to encode recursive data-types (e.g., lists). Two formulations of recursive types exist: iso-recursive and equi-recursive. The relative advantages of iso- and equi-recursion are well-studied when it comes to their impact on type-inference. However, the relative semantic expressiveness of the two formulations remains unclear so far.

This paper studies the semantic expressiveness of STLC with iso- and equi-recursive types, proving that these formulations are *equally expressive*. In fact, we prove that they are both as expressive as STLC with only term-level recursion. We phrase these equi-expressiveness results in terms of full abstraction of three canonical compilers between these three languages (STLC with iso-, with equi-recursive types and with term-level recursion). Our choice of languages allows us to study expressiveness when interacting over both a simply-typed and a recursively-typed interface. The three proofs all rely on a typed version of a proof technique called approximate backtranslation.

Together, our results show that there is no difference in semantic expressiveness between STLCs with iso- and equi-recursive types. In this paper, we focus on a simply-typed setting but we believe our results scale to more powerful type systems like System F.

CCS Concepts: • **Theory of computation** → **Lambda calculus**; *Type theory*; • **Software and its engineering** → **Recursion**.

Additional Key Words and Phrases: Fully-abstract compilation, Lambda Calculus, Recursive types, Iso-recursive types, Coinductive Equi-recursive types, Backtranslation

ACM Reference Format:

Marco Patrignani, Eric Mark Martin, and Dominique Devriese. 2021. On the Semantic Expressiveness of Recursive Types. *Proc. ACM Program. Lang.* 5, POPL, Article 21 (January 2021), 29 pages. <https://doi.org/10.1145/3434302>

21

To present notions more clearly, this paper uses syntax highlighting accessible to both colourblind and black & white readers [Patrignani 2020]. For a better experience, please print or view this in colour.¹

¹Specifically, we use a **blue, sans-serif** font for **STLC** with the **fix** operator, a **red, bold** font for **STLC** with **iso-recursive** types, and **pink, italic** font for **STLC** with **coinductive equi-recursive** types. Elements common to all languages are typeset in a **black, italic** font (to avoid repetition).

Authors' addresses: Marco Patrignani, Computer Science, Stanford University, USA, CISPA Helmholtz Center for Information Security, Saarbrücken, Germany, mp@cs.stanford.edu; Eric Mark Martin, Computer Science, Stanford University, USA, ericmarkmartin@cs.stanford.edu; Dominique Devriese, Computer Science, Vrije Universiteit Brussel, Brussels, Belgium, dominique.devriese@vub.be;

This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2021 Copyright held by the owner/author(s).

2475-1421/2021/1-ART21

<https://doi.org/10.1145/3434302>

1 INTRODUCTION

Recursive types were first proposed by Morris [1968] as a way to recover divergence from the untyped lambda calculus in a simply-typed lambda calculus. They also enable the definition of recursive data-types such as lists, trees, and Lisp S-expressions in typed languages.

Morris' original formulation was equi-recursive: a type $\mu\alpha. \tau$ was regarded as an infinite type and considered equal to its unfolding $\tau[\mu\alpha. \tau/\alpha]$. Subsequent formulations (e.g., Abadi and Fiore [1996]) use different type equality relations. In this paper we will work with λ_E^μ : a standard simply-typed lambda calculus with coinductive equi-recursive types [e.g., Cai et al. 2016].

Years after Morris' formulation of recursive types, a different one appeared [e.g., Gordon et al. 1979; Harper and Mitchell 1993], where the two types are not considered equal, but *isomorphic*: values can be converted from $\mu\alpha. \tau$ to $\tau[\mu\alpha. \tau/\alpha]$ and back using explicit **fold** and **unfold** annotations in terms. These annotations are used to guide typechecking, but they also have a significance at runtime: an explicit reduction step is needed to cancel them out: $\text{unfold}_{\mu\alpha. \tau} (\text{fold}_{\mu\alpha. \tau} v) \hookrightarrow v$. In this paper, we work with a standard iso-recursive calculus λ_I^μ .

The relation between these two formulations has been studied by Abadi and Fiore [1996] and Urzyczyn [1995] (the latter focusing on positive recursive types). Specifically, they show that any term typable in one formulation can also be typed in the other, possibly by adding extra **unfold** or **fold** annotations. Additionally, Abadi and Fiore prove that for types considered equal in the equi-recursive system, there exist coercion functions in the iso-recursive formulation that are mutually inverse in the (axiomatised) program logic. The isomorphism properties are proved in a logic for the iso-recursive language (which is only conjectured to be sound), and the authors do not even consider an operational semantics.

The relative semantic expressiveness of the two formulations, however, has remained yet unexplored. In principle, executions that are converging in the equi-recursive language may become diverging in the iso-recursive setting because of the extra fold-unfold reductions. Because of this, it is unclear whether the two formulations of recursive types produce equally expressive languages.

To study language expressiveness meaningfully, it is important to phrase the question properly. If we just consider programs that receive a natural number and return a boolean, then both languages will allow expressing the same set of algorithms, simply by their Turing completeness.

The question of expressiveness is more interesting if we consider programs that interact over a richer interface. Consider, for example, a term t from the simply-typed lambda calculus embedded into either calculus λ_I^μ or λ_E^μ . A much more interesting question is whether there are ways in which λ_E^μ contexts (i.e., larger programs) can interact with t that contexts in λ_I^μ cannot. The use of contexts in different languages interacting with a common term as a way of measuring language expressiveness has a long history [Felleisen 1991; Mitchell 1993], mostly in the study of process calculi [Parrow 2008]. In this setting, equal expressiveness of programming languages is sometimes argued for by proving the existence of a fully-abstract compiler from one language to the other [Gorla and Nestmann 2016]. Such a compiler translates contextually-equivalent terms in a source language (indicated as L_{src}) to contextually-equivalent terms in a target language (indicated as L_{trg}) [Abadi 1998; Patrignani et al. 2019]. That is, if contexts cannot distinguish two terms in L_{src} , they will also not be able to distinguish them after the compilation to L_{trg} .

Concretely, in this paper, we study the expressive power of λ_I^μ and λ_E^μ when interacting over two kinds of interfaces. The first is characterized by simply-typed lambda calculus types which do not mention recursive types themselves. We consider implementations of this interface in λ^{fx} (a simply typed lambda calculus with term-level recursion in the form of a primitive fixpoint operator), and embed them canonically into both λ_I^μ and λ_E^μ . We show that if two λ^{fx} terms cannot be distinguished by λ^{fx} contexts, then the same is true for both λ_I^μ and λ_E^μ contexts. Additionally, we consider STLC

types that contain recursive types themselves as interfaces. We take implementations of them in λ_1^μ and a canonical compiler into λ_E^μ . We show that this compiler is also fully abstract. These three fully-abstract compilation results establish the equi-expressiveness of λ_1^μ , λ_E^μ , and λ^{fx} contexts, interacting over simply-typed interfaces with and without recursive types.

Let us now argue why the choice of fully-abstract compilation as a measure of the relative expressiveness of programming languages is the right one in our setting. After all, several researchers have pointed out that the mere existence of a fully-abstract compilation is not in itself meaningful and only compilers that are sufficiently well-behaved should be considered [Gorla and Nestmann 2016; Parrow 2008]. The reason for this is that one can build a degenerate fully-abstract compiler that shows both languages having an equal amount (cardinality of) equivalence classes for terms. This would indicate that the languages are equally-expressive, but unfortunately this is also trivial to satisfy [Parrow 2008]. These degenerate examples, as such, clarify the necessity for well-behavedness of the compiler. However, we have not found a clear argument explaining why well-behaved fully-abstract compilation implies equi-expressiveness of languages, so here it is.

In our opinion (and we believe this point has not yet been made in the literature), the issue is that fully-abstract compilation results measure language expressiveness *not* by verifying that they can express the same *terms*, but that they can express the same *contexts*. Defining when a context in L_{src} is the same as a context in L_{trg} is hard, and therefore fully-abstract compilation simply requires that L_{trg} contexts can express the interaction of L_{src} contexts with any term that is shared between both languages. The role of the compiler, the translation from L_{src} to L_{trg} , is simply to obtain this common term against which expressiveness of contexts in both languages can be measured.

In other words, expressiveness of a programming language is only meaningful with respect to a certain interface and the role of the compiler is to map L_{src} implementations of this interface to L_{trg} implementations. In a sense, the L_{src} implementation of the interface should be seen as an expressiveness challenge for L_{src} contexts and the compiler translates it to the corresponding challenge in L_{trg} . As such, the compiler should be seen as part of the definition of equi-expressiveness and the well-behavedness requirement is there to make sure the L_{src} challenge is translated to “the same” challenge in L_{trg} . Fortunately, in this work we only rely on canonical compilers that provide the most intuitive translation for a term in our source languages into “the same” term in our target ones. Thus, we believe that in our setting using fully-abstract compilation is the right tool to measure the relative expressiveness of programming languages.

Proving full abstraction for a compiler is notoriously hard, particularly the preservation direction, i.e., showing that equivalent source terms get compiled to equivalent target terms. Informally, it requires showing that any behaviour (e.g., termination) of target program contexts can be replicated by source program contexts. Demonstrating such a claim is particularly complicated in our setting since λ_E^μ contexts have coinductive (and thus infinite) type equality derivations. To be able to prove fully-abstract compilation, we adopt the approximate backtranslation proof technique of Devriese et al. [2017]. This technique relies on two key components: a cross-language approximation relation between source and target terms (and source and target program contexts) and a backtranslation function from target to source program contexts. Intuitively, the approximation relation is used to tell when a source and a target term (or program context) equi-terminate; we use step-indexed logical relations to define this and rely on the step as the measure for the approximation. The backtranslation is a function that takes a target program context and produces a source program context that approximates the target one. This is particularly appropriate for backtranslating λ_E^μ program contexts, since we show that it is sufficient to approximate their coinductive derivations instead of replicating them precisely.

We construct three backtranslations: from λ_I^μ and λ_E^μ contexts respectively into λ^{fx} ones and from λ_E^μ contexts into λ_I^μ ones. We do so by defining a family of types for backtranslated terms that is not just indexed by the approximation level but also by the target type of the backtranslated term. To the best of our knowledge, this is a novel approach, since all existing work relies on a single type for backtranslated terms [Devriese et al. 2017; New et al. 2016]. For proving correctness of these backtranslations, we define a (step-)indexed logical relation to express when compiled and backtranslated terms approximate each other. While the logical relation is largely the same for the different compilers and backtranslations, differences in the language semantics impose that we treat backtranslated λ_I^μ terms differently from λ_E^μ .

To summarize, the key contribution of this paper is the proof that iso- and equi-recursive typing are equally expressive. This result is achieved via the following contributions (depicted in Figure 1).

- adapting the approximate backtranslation proof technique to operate on families of backtranslation types that are type-indexed on target types and compilers that do not rely on dynamic typechecks to attain fully-abstract compilation;
- proving that the compiler from λ^{fx} to λ_I^μ is fully abstract with an approximate backtranslation;
- proving that the compiler from λ^{fx} to λ_E^μ is fully abstract with an approximate backtranslation;
- proving that the compiler from λ_I^μ to λ_E^μ is fully abstract with an approximate backtranslation.

Note that technically, we can derive the compiler and backtranslation between λ^{fx} and λ_E^μ by composing the compilers and backtranslations through λ_I^μ . We present this result directly because it offers insights on proofs of fully-abstract compilation for languages with coinductive notions.

The remainder of this paper is organised as follows. We first formalise the languages we use (λ^{fx} , λ_I^μ and λ_E^μ) as well as the cross-language logical relations which express when two terms in those languages are semantically equivalent (Section 2). Next, we present fully-abstract compilation and describe our approximate backtranslation proof technique in detail (Section 3). Then we define the three compilers (from λ^{fx} to λ_I^μ , from λ^{fx} to λ_E^μ and from λ_I^μ to λ_E^μ) and prove that they are fully abstract using three approximate backtranslations (Section 4). Finally, we discuss related work and conclude (Sections 5 and 6).

For space constraints we omit some formalisation, auxiliary lemmas and proofs, which can be found in the online appendix [Patrignani et al. 2020].

2 LANGUAGES AND CROSS-LANGUAGE LOGICAL RELATIONS

This section presents the simply-typed lambda calculus (λ) and its extensions with a typed fixpoint operator (λ^{fx}), with iso-recursive types (λ_I^μ) and with coinductive equi-recursive types (λ_E^μ). We first define the syntax (Section 2.1), then the static semantics (Section 2.2) and then the operational semantics of these languages (Section 2.3). Finally, this section presents the cross-language logical relations used to reason about the expressiveness of terms in different languages (Section 2.4). Note

Fig. 1. Our contributions, visually. Full arrows indicate canonical embeddings $\llbracket \cdot \rrbracket$ while dotted ones are (approximate) backtranslations $\langle\langle \cdot \rangle\rangle$. Translations' superscripts indicate input languages while their subscripts indicate output languages.

that these logical relations are partial, the key addition needed to attain fully-abstract compilation is presented in Section 3.3 only after said addition is justified.

2.1 Syntax

All languages include standard terms (t) and values (v) from the simply-typed lambda calculus: lambda abstractions, applications, pairs, projections, tagged unions, case destructors, booleans, branching, unit and sequencing. Additionally, λ^{fx} has a **fix** operator providing general recursion, while λ_1^μ has **fold** and **unfold** annotations; λ_E^μ requires no additional syntactic construct. Regarding types, both λ_1^μ and λ_E^μ add recursive types according to the same syntax. In λ_1^μ and λ_E^μ , recursive types are syntactically constrained to be *contractive*. Note however that for simplicity of presentation we will indicate a type as τ and simply report the contractiveness constraints when meaningful. A recursive type $\mu\alpha. \tau$ is contractive if, the use of the recursion variable α in τ occurs under a type constructor such as $\rightarrow$ or $\times$ [MacQueen et al. 1984]. Non-contractive types (e.g., $\mu\alpha. \alpha$) are not inhabited by any value, so it is reasonable to elide them (Lemma 1). Moreover, they do not have an infinite unfolding and (without restrictions on the type equality relation) can be proven equivalent to any other type [Im et al. 2013], which is undesirable.

LEMMA 1 (NO VALUE HAS A NON-CONTRACTIVE TYPE). *if τ is non-contractive then $\nexists v. \emptyset \vdash v : \tau$.*

All languages have evaluation contexts ($\mathbb{E}$), which indicate where the next reduction will happen, and program contexts ($\mathbb{C}$), which are larger programs to link terms with.

$$\begin{aligned}
 v &::= \text{unit} \mid \text{true} \mid \text{false} \mid \lambda x : \tau. t \mid \langle v, v \rangle \mid \text{inl } v \mid \text{inr } v \mid \text{fold}_{\mu\alpha. \tau} v & \Gamma &::= \emptyset \mid \Gamma, x : \tau \\
 \tau, \sigma &::= \text{Unit} \mid \text{Bool} \mid \tau^s \rightarrow \tau^s \mid \tau^s \times \tau^s \mid \tau^s \uplus \tau^s \mid \mu\alpha. \tau \mid \mu\alpha. \tau & \tau^s &::= \alpha \mid \alpha \mid \tau \\
 t &::= \text{unit} \mid \text{true} \mid \text{false} \mid \lambda x : \tau. t \mid x \mid t \mid t.1 \mid t.2 \mid \langle t, t \rangle \mid \text{case } t \text{ of } \text{inl } x_1 \mapsto t \mid \text{inr } x_2 \mapsto t \\
 &\quad \mid \text{inl } t \mid \text{inr } t \mid \text{if } t \text{ then } t \text{ else } t \mid t; t \mid \text{fix}_{\tau \rightarrow \tau} t \mid \text{fold}_{\mu\alpha. \tau} t \mid \text{unfold}_{\mu\alpha. \tau} t \\
 \mathbb{E} &::= [\cdot] \mid \mathbb{E} \ t \mid v \ \mathbb{E} \mid \mathbb{E}.1 \mid \mathbb{E}.2 \mid \langle \mathbb{E}, t \rangle \mid \langle v, \mathbb{E} \rangle \mid \text{case } \mathbb{E} \text{ of } \text{inl } x_1 \mapsto t_1 \mid \text{inr } x_2 \mapsto t_2 \\
 &\quad \mid \text{inl } \mathbb{E} \mid \text{inr } \mathbb{E} \mid \mathbb{E}; t \mid \text{if } \mathbb{E} \text{ then } t \text{ else } t \mid \text{fix}_{\tau \rightarrow \tau} \mathbb{E} \mid \text{fold}_{\mu\alpha. \tau} \mathbb{E} \mid \text{unfold}_{\mu\alpha. \tau} \mathbb{E} \\
 \mathbb{C} &::= [\cdot] \mid \lambda x : \tau. \mathbb{C} \mid \mathbb{C} \ t \mid t \ \mathbb{C} \mid \mathbb{C}.1 \mid \mathbb{C}.2 \mid \langle \mathbb{C}, t \rangle \mid \langle t, \mathbb{C} \rangle \mid \text{case } \mathbb{C} \text{ of } \text{inl } x_1 \mapsto t \mid \text{inr } x_2 \mapsto t \\
 &\quad \mid \text{case } t \text{ of } \text{inl } x_1 \mapsto \mathbb{C} \mid \text{inr } x_2 \mapsto \mathbb{C} \mid \text{case } t \text{ of } \text{inl } x_1 \mapsto t \mid \text{inr } x_2 \mapsto \mathbb{C} \mid \text{inl } \mathbb{C} \\
 &\quad \mid \text{inr } \mathbb{C} \mid \mathbb{C}; t \mid t; \mathbb{C} \mid \text{if } \mathbb{C} \text{ then } t \text{ else } t \mid \text{if } t \text{ then } \mathbb{C} \text{ else } t \mid \text{if } t \text{ then } t \text{ else } \mathbb{C} \\
 &\quad \mid \text{fix}_{\tau \rightarrow \tau} \mathbb{C} \mid \text{fold}_{\mu\alpha. \tau} \mathbb{C} \mid \text{unfold}_{\mu\alpha. \tau} \mathbb{C}
 \end{aligned}$$

2.2 Static Semantics

This section presents the (fairly standard) static semantics of our languages, we delay discussing alternative formulations of equi-recursive types to Section 5. The static semantics for terms follows the canonical judgement $\Gamma \vdash t : \tau$, which attributes type τ to term t under environment Γ and occasionally relies on function $\text{ftv}(\tau)$, which returns the free type variables of τ . The only difference in the typing rules regards **fold/unfold** terms (Rules λ_1^μ -Type-fold and λ_1^μ -Type-unfold) and the introduction of the type equality ($\cong$ in Rule λ_E^μ -Type-eq).

$\Gamma \vdash t : \tau$			
(Type-var) $x : \tau \in \Gamma$ $\Gamma \vdash x : \tau$	(Type-unit) $\Gamma \vdash \text{unit} : \text{Unit}$	(Type-true) $\Gamma \vdash \text{true} : \text{Bool}$	(Type-false) $\Gamma \vdash \text{false} : \text{Bool}$
(Type-lam) $\Gamma, x : \tau \vdash t : \tau'$ $\Gamma \vdash \lambda x : \tau. t : \tau \rightarrow \tau'$	(Type-pair) $\text{ftv}(\tau) = \emptyset$ $\Gamma \vdash t : \tau \quad \Gamma \vdash t' : \tau'$ $\Gamma \vdash \langle t, t' \rangle : \tau \times \tau'$	(Type-inl) $\Gamma \vdash t : \tau$ $\Gamma \vdash \text{inl } t : \tau \uplus \tau'$	

$$\begin{array}{c}
\text{(Type-inr)} \quad \frac{\Gamma \vdash t : \tau'}{\Gamma \vdash \text{inr } t : \tau \uplus \tau'} \quad \text{(Type-app)} \quad \frac{\Gamma \vdash t : \tau' \rightarrow \tau \quad \Gamma \vdash t' : \tau'}{\Gamma \vdash t t' : \tau} \quad \text{(Type-p1)} \quad \frac{\Gamma \vdash t : \tau \times \tau'}{\Gamma \vdash t.1 : \tau} \quad \text{(Type-p2)} \quad \frac{\Gamma \vdash t : \tau' \times \tau}{\Gamma \vdash t.2 : \tau} \\
\text{(Type-case)} \quad \frac{\Gamma \vdash t : \tau' \uplus \tau'' \quad \Gamma, x_1 : \tau' \vdash t' : \tau \quad \Gamma, x_2 : \tau'' \vdash t'' : \tau}{\Gamma \vdash \text{case } t \text{ of } \text{inl } x_1 \mapsto t' \mid \text{inr } x_2 \mapsto t'' : \tau} \quad \text{(Type-seq)} \quad \frac{\Gamma \vdash t : \text{Unit} \quad \Gamma \vdash t' : \tau}{\Gamma \vdash t; t' : \tau} \\
\text{(Type-if)} \quad \frac{\Gamma \vdash t : \text{Bool} \quad \Gamma \vdash t' : \tau \quad \Gamma \vdash t'' : \tau}{\Gamma \vdash \text{if } t \text{ then } t' \text{ else } t'' : \tau} \quad \text{(\lambda}^{\text{fx}}\text{-Type-fix)} \quad \frac{\Gamma \vdash t : (\tau_1 \rightarrow \tau_2) \rightarrow \tau_1 \rightarrow \tau_2}{\Gamma \vdash \text{fix}_{\tau_1 \rightarrow \tau_2} t : \tau_1 \rightarrow \tau_2} \\
\text{(\lambda}_1^\mu\text{-Type-fold)} \quad \frac{\Gamma \vdash t : \tau[\mu\alpha. \tau/\alpha]}{\Gamma \vdash \text{fold}_{\mu\alpha. \tau} t : \mu\alpha. \tau} \quad \text{(\lambda}_1^\mu\text{-Type-unfold)} \quad \frac{\Gamma \vdash t : \mu\alpha. \tau}{\Gamma \vdash \text{unfold}_{\mu\alpha. \tau} t : \tau[\mu\alpha. \tau/\alpha]} \quad \text{(\lambda}_E^\mu\text{-Type-eq)} \quad \frac{\Gamma \vdash t : \mu\alpha. \tau \quad \mu\alpha. \tau \doteq \sigma}{\Gamma \vdash t : \sigma}
\end{array}$$

Program contexts have an important role in fully-abstract compilation. They follow the usual typing judgement ($\mathbb{C} \vdash \Gamma, \tau \rightarrow \Gamma', \tau'$), i.e., program context $\mathbb{C}$ is well typed with a hole of type τ that use free variables in Γ , and overall $\mathbb{C}$ returns a term of type τ' and uses variables in Γ' . These typing rules are unsurprising, so we omit them for space constraints.

We use the same coinductive type equality relation of Cai et al. [2016], with a cosmetic difference only. Two types are equal if they are the same base type ι or variable (Rules $\doteq\text{-prim}$ and $\doteq\text{-var}$). If the types are composed of two types, the connectors must be the same and each sub-type must be equivalent (Rule $\doteq\text{-bin}$). If the left type starts with a μ (or if that does not but the right one does), then we unfold the type for checking the equality (Rules $\doteq\text{-}\mu_l$ and $\doteq\text{-}\mu_r$). Note that these last two rules are defined in an asymmetric fashion to make equality derivation deterministic. Finally, we make explicit the rules for reflexivity, symmetry and transitivity (Rules $\doteq\text{-refl}$ to $\doteq\text{-trans}$) whose derivations we have proved from the other rules.

$$\begin{array}{c}
\boxed{\tau \doteq \tau'} \\
\hline
\text{(\doteq-prim)} \quad \frac{\iota = \text{Unit} \vee \iota = \text{Bool}}{\iota \doteq \iota} \quad \text{(\doteq-var)} \quad \frac{}{\alpha \doteq \alpha} \quad \text{(\doteq-bin)} \quad \frac{\star \in \{\rightarrow, \times, \uplus\} \quad \tau_1 \doteq \sigma_1 \quad \tau_2 \doteq \sigma_2}{\tau_1 \star \tau_2 \doteq \sigma_1 \star \sigma_2} \quad \text{(\doteq-}\mu_l\text{)} \quad \frac{\tau[\mu\alpha. \tau/\alpha] \doteq \sigma \quad \tau \text{ contractive in } \alpha}{\mu\alpha. \tau \doteq \sigma} \\
\text{(\doteq-}\mu_r\text{)} \quad \frac{\text{lmc}(\tau) = 0 \quad \tau \doteq \sigma[\mu\alpha. \sigma/\alpha] \quad \sigma \text{ contractive in } \alpha}{\tau \doteq \mu\alpha. \sigma} \quad \text{(\doteq-refl)} \quad \frac{}{\tau \doteq \tau} \quad \text{(\doteq-symm)} \quad \frac{\sigma \doteq \tau}{\tau \doteq \sigma} \quad \text{(\doteq-trans)} \quad \frac{\tau \doteq \sigma \quad \sigma \doteq \tau'}{\tau \doteq \tau'}
\end{array}$$

To prove results about this equality relation, we will often induct on the “leading-mu-count” (lmc) measure. Intuitively, that measure counts the amount of μ s that a λ_E^μ type has before a different connector is found. This is almost the same as the number of times a type can be unfolded before it is no longer recursive at the top level (e.g. $\text{lmc}(\text{Unit}) = 0$, $\text{lmc}(\mu\alpha. \alpha \uplus \text{Unit}) = 1$). Non-contractive types such as $\mu\alpha. \alpha$, however, create problems here, for they always unfold into another top level recursive type. This motivates our restriction to contractive types only: a contractive type τ can be unfolded exactly $\text{lmc}(\tau)$ times. This restriction is harmless, since non-contractive recursive types are not inhabited by any value (Lemma 1).

2.3 Dynamic Semantics

All our languages are given a contextual, call-by-value, operational semantics. We highlight primitive reductions as $\hookrightarrow_p$ and non-primitive ones as $\hookrightarrow$. We indicate the capture-avoiding substitution

of variable (or type variable) x in t with value (or type) v as $t[v/x]$. Note that since λ_E^μ has no peculiar syntactic construct, it also has no specific reduction rule.

$$\begin{array}{c}
 \hline
 \boxed{t \hookrightarrow t' \quad \text{and} \quad t \hookrightarrow_p t'} \\
 \hline
 \begin{array}{c}
 \text{(Eval-ctx)} \\
 \frac{t \hookrightarrow_p t'}{\mathbb{E}[t] \hookrightarrow \mathbb{E}[t']} \\
 \text{(Eval-beta)} \\
 \frac{(\lambda x : \tau. t) \ v \hookrightarrow_p t[v/x]}{\text{(Eval-inl)}} \\
 \text{(Eval-pi)} \\
 \frac{i \in 1..2}{\langle v_1, v_2 \rangle . i \hookrightarrow_p v_i} \\
 \text{(Eval-seq)} \\
 \frac{\text{unit}; t \hookrightarrow_p t}{\text{(Eval-inr)}} \\
 \frac{\text{case inl } v \text{ of } \left\{ \begin{array}{l} \text{inl } x_1 \mapsto t \\ \text{inr } x_2 \mapsto t' \end{array} \right. \hookrightarrow_p t[v/x_1]}{\text{(Eval-if)}} \\
 \frac{v = \text{true} \vee \text{false}}{\text{if } v \text{ then } t_{\text{true}} \text{ else } t_{\text{false}} \hookrightarrow_p t_v} \\
 \frac{\text{case inr } v \text{ of } \left\{ \begin{array}{l} \text{inl } x_1 \mapsto t \\ \text{inr } x_2 \mapsto t' \end{array} \right. \hookrightarrow_p t'[v/x_2]}{\text{(\lambda^{fx}.\text{Eval-fix})}} \\
 \frac{\text{fix}_{\tau \rightarrow \tau} \ (\lambda x : \tau. t) \hookrightarrow_p t \ [\text{fix}_{\tau \rightarrow \tau} \ \lambda x : \tau. t/x]}{\text{(\lambda_I^\mu - Eval-fold)}} \\
 \hline
 \text{unfold}_{\mu\alpha.\tau} \ (\text{fold}_{\mu\alpha.\tau} \ v) \hookrightarrow_p v
 \end{array}
 \end{array}$$

2.4 Logical Relations Between Our Languages

As mentioned in Section 1, we need cross-language relations that indicate when related source and target terms approximate each other. Intuitively, one such relation is needed by each one of the compilers we define later. Thus, we need to define three logical relations: A one between λ^{fx} and λ_I^μ , which we dub $LR_{\mu I}^{\text{fx}}$; B one between λ^{fx} and λ_E^μ , which we dub $LR_{\mu E}^{\text{fx}}$; C one between λ_I^μ and λ_E^μ , which we dub $LR_{\mu E}^{\mu I}$. They are all indexed by (a step and then by) the source type, so logical relations (A) and (B) look the same. For brevity we present only one of them. Additionally, given that λ_I^μ has the same types of λ^{fx} plus recursive types, we only show that case for logical relation (C). Ours are Kripke, step-indexed logical relations that are based on those of [Devriese et al. \[2017\]](#); [Hur and Dreyer \[2011\]](#). The step-indexing is not inherently needed for relations (A) and (B), which could be defined just by induction on λ^{fx} types (since they do not include recursive types). However, all of our relations are step-indexed anyway because the steps also determine for how many steps one term should approximate the other.

Before presenting the details, note that the relations we show here are *not* complete. Specifically they only talk about the terms needed to conclude reflection of fully-abstract compilation but not preservation (admittedly, the most interesting part). Completing the logical relations relies on technical insights regarding the backtranslations, so we do this later in Section 3.3. The goal of this section is to provide an understanding of what it means for two terms to approximate each other.

All three relations rely on the same notion of very simple Kripke worlds W (Fig. 2). Worlds consist of just a step-index k that is accessed via function $\text{lev}(W)$. The $\triangleright$ modality and future world relation $\sqsupseteq$ express that future worlds allow programs to take fewer reduction steps. We define two different observation relations, one for each direction of the approximations we are interested in: $O(W)_{\lesssim}$ and $O(W)_{\gtrsim}$ while $O(W)_{\approx}$ indicates the intersection of those approximations. Both these relation use notation $t \hookrightarrow^n v$, which indicates that term t reduces to value v in n steps or less. The former defines that a source term approximates a target term if termination of the first in $\text{lev}(W)$ steps or less implies termination of the second (in any number of steps). The latter requires the reverse. All of our logical relations will be defined in terms of either $O(W)_{\lesssim}$ or $O(W)_{\gtrsim}$. For definitions and lemmas or theorems that apply for both instantiations, we use the symbol ∇ as a metavariable that can be instantiated to either $\lesssim$ or $\gtrsim$.

Note that our logical relations (Figure 3) are not indexed by source types, but by *pseudo-types* $\hat{t}$. Pseudo-types contain all the constructs of source types, plus an additional type which we indicate for now as $EmulT$. This type is not a source type; it is needed because of the approximate backtranslation, so we defer explaining its details until Section 3.3.

Function $\text{repEmul}^{\text{fI}}(\cdot)$ converts a pseudo-type to an actual source type by replacing all occurrences of $EmulT$ with a concrete source type.² We will sometimes silently use a normal source type where a pseudo-type is expected; this makes sense since the syntax for the latter is a superset of the former. Function $\text{fxToIs}(\cdot)$ converts a source pseudo-type into its target-level correspondent; this is needed because unlike the previous work of Devriese et al. [2017], all of our target languages are typed. The formal details of both these functions are deferred until $EmulT$ is defined (Section 3.3). Finally, function $\text{ofType}^{\text{fI}}(\cdot)$ checks that terms have the correct form according to the rules of syntactic typing (Section 2.2). Function $\text{ofType}^{\text{IE}}(\cdot)$ does the analogous syntactic typecheck but for terms of λ_1^μ and λ_E^μ .

The value relation $\mathcal{V} \llbracket \hat{t} \rrbracket_\gamma$ is defined inductively on source pseudo-types and it is quite standard. *Unit* and *Bool* values are related in any world so long as they are the same value. Function values are related if they are well-typed, if both are lambdas, and if substituting related values in the bodies yields related terms in any strictly-future world. Pair values are related if both are pairs and each projection is related in strictly-future worlds and sum values are related if they have the same tag (*inl* or *inr*) and the tagged values are related in strictly-future worlds. Finally, the value relation for recursive types used by $LR_{\mu E}^{\mu I}$ is not defined on strictly-future worlds because in an equi-recursive language, values of recursive type can be inspected without consuming a step. However, this does not compromise well-foundedness of the relation because our recursive types $\mu\alpha. \tau$ are contractive, so the recursion variable α in τ must occur under a type constructor such as $\rightarrow$ and the relation for these constructors recurses only at strictly-future worlds.

The value, evaluation context and term relations are defined by mutual recursion, using a technique called biorthogonality (see, e.g., [Benton and Hur 2009]). Evaluation contexts $\mathcal{K} \llbracket \hat{t} \rrbracket_\gamma$ are related in a world if plugging in related values in any future world yields terms that are related according to the observation relation of the world. Similarly, terms are related $\mathcal{E} \llbracket \hat{t} \rrbracket_\gamma$ if plugging the terms in related evaluation contexts yields terms related according to the observation relation of the world. Relation $\mathcal{G} \llbracket \hat{t} \rrbracket_\gamma$ relates substitutions; this simply requires that substitutions for all variables in the context are for related values.

We indicate open terms to be logically related according to the three relations as follows:

$$LR_{\mu I}^{\text{fx}} : \hat{t} \vdash t \nabla_n t : \hat{t} \quad LR_{\mu E}^{\text{fx}} : \hat{t} \vdash t \nabla_n t : \hat{t} \quad LR_{\mu E}^{\mu I} : \hat{t} \vdash t \nabla_n t : \hat{t}$$

An open source term is related up to n steps at pseudo-type $\hat{t}$ in pseudo-context $\hat{t}$ to a target open term if both are well-typed and closing both terms with substitutions related in $\hat{t}$ produces terms related at $\hat{t}$ in any world that has at least n steps. If terms are related for any number of steps,

² As a convention, superscripts of these auxiliary functions indicate the initials of the two languages involved.

$$\begin{aligned} W &\stackrel{\text{def}}{=} n \in \mathbb{N} \quad \text{lev}(n) = n \quad \triangleright(0) = 0 \quad \triangleright(n+1) = n \\ W \sqsubseteq W' &= \text{lev}(W) \leq \text{lev}(W') \quad W \sqsupset W' = \text{lev}(W) < \text{lev}(W') \\ O(W)_{\leq} &\stackrel{\text{def}}{=} \{ (t, t) \mid \text{if } \text{lev}(W) > n \text{ and } t \hookrightarrow^n v \text{ then } \exists k. t \hookrightarrow^k v \} \\ O(W)_{\geq} &\stackrel{\text{def}}{=} \{ (t, t) \mid \text{if } \text{lev}(W) > n \text{ and } t \hookrightarrow^n v \text{ then } \exists k. t \hookrightarrow^k v \} \\ O(W)_{\approx} &\stackrel{\text{def}}{=} O(W)_{\leq} \cap O(W)_{\geq} \end{aligned}$$

Fig. 2. Worlds, observations and related technicalities. These are typeset for the relation between λ^{fx} and λ_1^μ but the other ones do not change.

$$\begin{aligned}
\hat{\tau} &::= \text{Unit} \mid \text{Bool} \mid \hat{\tau} \rightarrow \hat{\tau} \mid \hat{\tau} \times \hat{\tau} \mid \hat{\tau} \uplus \hat{\tau} \mid \text{EmulT} \text{ (to be defined in Section 3.3)} \\
\text{oftype}^{\text{fI}}(\hat{\tau}) &\stackrel{\text{def}}{=} \{(\mathbf{v}, \mathbf{v}) \mid \mathbf{v} \in \text{oftype}(\hat{\tau}) \text{ and } \mathbf{v} \in \text{oftype}(\text{fxToIs}(\hat{\tau}))\} \\
\text{oftype}(\hat{\tau}) &\stackrel{\text{def}}{=} \{\mathbf{v} \mid \emptyset \vdash \mathbf{v} : \text{repEmul}^{\text{fI}}(\hat{\tau})\} \quad \text{oftype}(\tau) \stackrel{\text{def}}{=} \{\mathbf{v} \mid \emptyset \vdash \mathbf{v} : \tau\} \\
\text{repEmul}^{\text{fI}}(\cdot) : \hat{\tau} &\rightarrow \tau \text{ (see Section 3.3)} \quad \text{fxToIs}(\cdot) : \hat{\tau} \rightarrow \tau \text{ (see Section 3.3)}
\end{aligned}$$

$$\begin{aligned}
\triangleright R &\stackrel{\text{def}}{=} \{(W, \mathbf{v}, \mathbf{v}) \mid \text{if } \text{lev}(W) > 0 \text{ then } (\triangleright(W), \mathbf{v}, \mathbf{v}) \in R\} \\
\mathcal{V}[\![\text{Unit}]\!]_{\nabla} &\stackrel{\text{def}}{=} \{(W, \mathbf{v}, \mathbf{v}) \mid \mathbf{v} = \text{unit} \text{ and } \mathbf{v} = \text{unit}\} \\
\mathcal{V}[\![\text{Bool}]\!]_{\nabla} &\stackrel{\text{def}}{=} \{(W, \mathbf{v}, \mathbf{v}) \mid (\mathbf{v} = \text{true} \text{ and } \mathbf{v} = \text{true}) \text{ or } (\mathbf{v} = \text{false} \text{ and } \mathbf{v} = \text{false})\} \\
\mathcal{V}[\![\hat{\tau} \rightarrow \hat{\tau}']\!]_{\nabla} &\stackrel{\text{def}}{=} \left\{ (W, \mathbf{v}, \mathbf{v}) \left| \begin{array}{l} (\mathbf{v}, \mathbf{v}) \in \text{oftype}^{\text{fI}}(\hat{\tau} \rightarrow \hat{\tau}') \text{ and} \\ \exists \mathbf{t}, \mathbf{t}. \mathbf{v} = \lambda \mathbf{x} : \text{repEmul}^{\text{fI}}(\hat{\tau}). \mathbf{t}, \mathbf{v} = \lambda \mathbf{x} : \text{fxToIs}(\hat{\tau}). \mathbf{t} \text{ and} \\ \forall W', \mathbf{v}', \mathbf{v}'. \text{ if } W' \sqsupseteq W \text{ and } (W', \mathbf{v}', \mathbf{v}') \in \mathcal{V}[\![\hat{\tau}]\!]_{\nabla} \text{ then} \\ (W', \mathbf{t}[\mathbf{v}'/\mathbf{x}], \mathbf{t}[\mathbf{v}'/\mathbf{x}]) \in \mathcal{E}[\![\hat{\tau}']\!]_{\nabla} \end{array} \right. \right\} \\
\mathcal{V}[\![\hat{\tau} \times \hat{\tau}']\!]_{\nabla} &\stackrel{\text{def}}{=} \left\{ (W, \mathbf{v}, \mathbf{v}) \left| \begin{array}{l} (\mathbf{v}, \mathbf{v}) \in \text{oftype}^{\text{fI}}(\hat{\tau} \times \hat{\tau}') \text{ and} \\ \exists \mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_1, \mathbf{v}_2. \mathbf{v} = \langle \mathbf{v}_1, \mathbf{v}_2 \rangle, \mathbf{v} = \langle \mathbf{v}_1, \mathbf{v}_2 \rangle \text{ and} \\ (W, \mathbf{v}_1, \mathbf{v}_1) \in \triangleright \mathcal{V}[\![\hat{\tau}]\!]_{\nabla} \text{ and } (W, \mathbf{v}_2, \mathbf{v}_2) \in \triangleright \mathcal{V}[\![\hat{\tau}']\!]_{\nabla} \end{array} \right. \right\} \\
\mathcal{V}[\![\hat{\tau} \uplus \hat{\tau}']\!]_{\nabla} &\stackrel{\text{def}}{=} \left\{ (W, \mathbf{v}, \mathbf{v}) \left| \begin{array}{l} (\mathbf{v}, \mathbf{v}) \in \text{oftype}^{\text{fI}}(\hat{\tau} \uplus \hat{\tau}') \text{ and either} \\ \exists \mathbf{v}', \mathbf{v}'. (W, \mathbf{v}', \mathbf{v}') \in \triangleright \mathcal{V}[\![\hat{\tau}]\!]_{\nabla} \text{ and } \mathbf{v} = \text{inl } \mathbf{v}', \mathbf{v} = \text{inl } \mathbf{v}' \text{ or} \\ \exists \mathbf{v}', \mathbf{v}'. (W, \mathbf{v}', \mathbf{v}') \in \triangleright \mathcal{V}[\![\hat{\tau}']\!]_{\nabla} \text{ and } \mathbf{v} = \text{inr } \mathbf{v}', \mathbf{v} = \text{inr } \mathbf{v}' \end{array} \right. \right\} \\
\mathcal{V}[\![\text{EmulT}]\!]_{\nabla} &\stackrel{\text{def}}{=} \text{to be defined in Section 3.3} \\
\mathcal{K}[\![\hat{\tau}]\!]_{\nabla} &\stackrel{\text{def}}{=} \left\{ (W, \mathbb{E}, \mathbb{E}) \left| \begin{array}{l} \forall W', \mathbf{v}, \mathbf{v}. \text{ if } W' \sqsupseteq W \text{ and } (W', \mathbf{v}, \mathbf{v}) \in \mathcal{V}[\![\hat{\tau}]\!]_{\nabla} \text{ then} \\ (\mathbb{E}[\mathbf{v}], \mathbb{E}[\mathbf{v}]) \in O(W')_{\nabla} \end{array} \right. \right\} \\
\mathcal{E}[\![\hat{\tau}]\!]_{\nabla} &\stackrel{\text{def}}{=} \{(W, \mathbf{t}, \mathbf{t}) \mid \forall \mathbb{E}, \mathbb{E}. \text{ if } (W, \mathbb{E}, \mathbb{E}) \in \mathcal{K}[\![\hat{\tau}]\!]_{\nabla} \text{ then } (\mathbb{E}[\mathbf{t}], \mathbb{E}[\mathbf{t}]) \in O(W)_{\nabla}\} \\
\mathcal{G}[\![\emptyset]\!]_{\nabla} &\stackrel{\text{def}}{=} \{(W, \emptyset, \emptyset)\} \\
\mathcal{G}[\![\hat{\tau}, \mathbf{x} : \hat{\tau}]\!]_{\nabla} &\stackrel{\text{def}}{=} \{(W, \mathbf{y}[\mathbf{v}/\mathbf{x}], \mathbf{y}[\mathbf{v}/\mathbf{x}]) \mid (W, \mathbf{y}, \mathbf{y}) \in \mathcal{G}[\![\hat{\tau}]\!]_{\nabla} \text{ and } (W, \mathbf{v}, \mathbf{v}) \in \mathcal{V}[\![\hat{\tau}]\!]_{\nabla}\}
\end{aligned}$$

$$\mathcal{V}[\![\mu\hat{\alpha}. \tau]\!]_{\nabla} \stackrel{\text{def}}{=} \left\{ (W, \mathbf{v}, \mathbf{v}) \left| \begin{array}{l} (\mathbf{v}, \mathbf{v}) \in \text{oftype}^{\text{IE}}(\mu\hat{\alpha}. \tau) \text{ and} \\ \exists \mathbf{v}'. (W, \mathbf{v}', \mathbf{v}) \in \mathcal{V}[\![\tau[\mu\hat{\alpha}. \tau/\alpha]]\!]_{\nabla} \text{ and } \mathbf{v} = \text{fold}_{\mu\hat{\alpha}. \tau} \mathbf{v}' \end{array} \right. \right\}$$

Fig. 3. Part of the cross-language logical relation we rely on (classical bits) and its auxiliary functions.

we simply omit the n index and write $\hat{\Gamma} \vdash \mathbf{t} \nabla \mathbf{t} : \hat{\tau}$. Since we have to also relate program contexts across languages, we define what it means for them to be related as follows:

$$LR_{\mu\mathbf{I}}^{\text{fx}} : \vdash \mathbb{C} \nabla \mathbb{C} : \hat{\Gamma}, \hat{\tau} \rightarrow \hat{\Gamma}', \hat{\tau}' \quad LR_{\mu\mathbf{E}}^{\text{fx}} : \vdash \mathbb{C} \nabla \mathbb{C} : \hat{\Gamma}, \hat{\tau} \rightarrow \hat{\Gamma}', \hat{\tau}' \quad LR_{\mu\mathbf{E}}^{\mu\mathbf{I}} : \vdash \mathbb{C} \nabla \mathbb{C} : \hat{\Gamma}, \hat{\tau} \rightarrow \hat{\Gamma}', \hat{\tau}'$$

Program contexts are related if they are well-typed and if plugging terms related at the pseudo-type of the hole ($\hat{\tau}$) in each of them produces terms related at the pseudo-type of the result ($\hat{\tau}'$).³

³ The interested reader will find the formalisation of these definitions in the online appendix [Patrignani et al. 2020].

All our logical relations are constructed so that for related terms, termination of one term implies termination of the other according to the direction of the approximation ($\lesssim$ or $\gtrsim$) (Lemma 2). We define termination of a term t as reduction to a value in some steps: $t \Downarrow \stackrel{\text{def}}{=} \exists n, v. t \hookrightarrow^n v$.

LEMMA 2 (ADEQUACY FOR $\approx$).

if $\emptyset \vdash t \lesssim_n t : \tau$ and $t \hookrightarrow^m v$ with $n \geq m$ then $t \Downarrow$ if $\emptyset \vdash t \gtrsim_n t : \tau$ and $t \hookrightarrow^m v$ with $n \geq m$ then $t \Downarrow$

3 FULLY-ABSTRACT COMPILATION AND APPROXIMATE BACKTRANSLATIONS

This section provides an overview of fully-abstract compilation and of the approximate backtranslation proof technique that we use (Section 3.1). The approximate backtranslation requires defining the backtranslation type, i.e., the type that represents backtranslated values (Section 3.2). This type provides the insights needed to complete the definitions of our logical relations and to understand how to reason about backtranslated terms cross-languages (Section 3.3).

3.1 A Primer on Fully-Abstract Compilation and Approximate Backtranslations

A compiler is fully abstract if it preserves and reflects contextual equivalence between source and target language [Abadi 1998]. Many compiler passes have been proven to satisfy this criterion [Ahmed and Blume 2008, 2011; Devriese et al. 2017; Fournet et al. 2013; New et al. 2016; Patrignani et al. 2015; Skorstengaard et al. 2019; Van Strydonck et al. 2019], we refer the interested reader to the survey of Patrignani et al. [2019].

Two programs are contextually equivalent if they produce the same behaviour no matter the larger program (i.e., program context) they interact with [Plotkin 1977]. As commonly done, we define “producing the same behaviour” as equi-termination (one terminates iff the other does). We use a complete formulation of contextual equivalence for typed programs, which enforces that contexts are well-typed and their types match that of the terms considered.

Definition 1 (Contextual Equivalence).

$$\Gamma \vdash t_1 \simeq_{\text{ctx}} t_2 : \tau \stackrel{\text{def}}{=} \Gamma \vdash t_1 : \tau \text{ and } \Gamma \vdash t_2 : \tau \text{ and } \forall \mathbb{C}. \mathbb{C} : \Gamma, \tau \rightarrow \emptyset, \tau'. \mathbb{C}[t_1] \Downarrow \iff \mathbb{C}[t_2] \Downarrow$$

Quantifying over all contexts in Definition 1 ensures that contextually-equivalent terms do not just equi-terminate, but that any value the context can obtain from them is indistinguishable.

For a compiler $\llbracket \cdot \rrbracket$ from language L_{src} to L_{trg} , we define full abstraction as follows:

Definition 2 (Fully-abstract compilation).

$$\vdash \llbracket \cdot \rrbracket : FA \stackrel{\text{def}}{=} \forall t_1, t_2 \in L_{\text{src}}. \emptyset \vdash t_1 \simeq_{\text{ctx}} t_2 : \tau \iff \emptyset \vdash \llbracket t_1 \rrbracket \simeq_{\text{ctx}} \llbracket t_2 \rrbracket : \llbracket \tau \rrbracket$$

For simplicity, we instantiate Definition 2 for closed terms only (i.e., well-typed under empty environments). Opening the environment to a non-empty set of term variables is straightforward and therefore omitted [Devriese et al. 2017].

3.1.1 Proving Fully-Abstract Compilation: Reflection (or, the Easy Part). The reflection part of fully-abstract compilation requires that the compiler produces equivalent target programs only if their source counterparts were equivalent. Contrapositively, inequivalent source programs must be compiled to inequivalent target program. This proof can often be derived as a corollary of standard compiler correctness (i.e., refinement) [Patrignani et al. 2019].

As mentioned, we prove the reflection direction by relying on the cross-language logical relations. Our logical relations are compiler-agnostic—they simply state when terms approximate each others (recall that $\approx$ is the intersection of both approximations $\lesssim$ and $\gtrsim$). However, we use them to show that any term (and program context) is related to its compilation. With this fact, by relying on the adequacy of logical relations (Lemma 2), we know that related terms equi-terminate. Thus, we can apply the reasoning depicted in Figure 4 (left) to conclude this part of fully-abstract compilation.

Fig. 4. Diagram breakdown of the reflection (left) and preservation (right) proofs of fully-abstract compilation.

3.1.2 Proving Fully-Abstract Compilation: Preservation (or, the Hard Part). Fully-abstract compilation proofs are notorious and their complexity resides in the *preservation* direction. That is, starting from contextually-equivalent programs in the source, prove that their compiled counterparts are contextually-equivalent in the target. For our three fully-abstract compilation results we rely on the approximate backtranslation proof technique [Devriese et al. 2017], depicted in Figure 4 (right).

We rely on both directions of the cross-language approximation relating terms for this proof. Recall that $t \gtrsim_n t$ is used to know that if t terminates in n steps in the target, then t also terminates (in arbitrary steps) in the source. The converse, $t \lesssim_n t$ is used to know that if t terminates in n steps in the source, then t also terminates (again in arbitrary steps) in the target. We start with source term t approximating (in both directions) its compilation $\llbracket t \rrbracket$. Then, to prove target contextual equivalence (the $\approx$ -decorated equivalence), we start by assuming that a target context $\llbracket t_1 \rrbracket$ linked with $\llbracket t_1 \rrbracket$ terminates in some steps ($\Downarrow_n$). Eventually, we need to show that the same target context linked with $\llbracket t_2 \rrbracket$ also terminates in any steps ($\Downarrow$). This is the $\approx$ -decorated implication, the reverse direction holds by symmetry. To progress, we construct a *backtranslation* $\langle\langle \cdot \rangle\rangle_n$, i.e., a function that takes a target context $\llbracket t \rrbracket$ and returns a source context that approximates $\llbracket t \rrbracket$ in both directions. With the backtranslation and this direction of the approximation $\gtrsim_n$, we prove implication (1): the backtranslated context $\langle\langle \llbracket t \rrbracket \rangle\rangle_n$ linked with t_1 terminates in the source. At this point, the assumption of source contextual equivalence yields implication (2): the same backtranslated context $\langle\langle \llbracket t \rrbracket \rangle\rangle_n$ linked with t_2 also terminates. Now we rely on the another direction of the approximation between the target context and its backtranslation (as well as between source terms and their compilation): $\lesssim$. This other approximation lets us conclude implication (3): the original target context $\llbracket t \rrbracket$ linked with $\llbracket t_2 \rrbracket$ terminates in the target. This is what we prove for a compiler to be fully abstract.

3.2 A Family of Backtranslation Types

Backtranslated contexts must be valid source contexts, i.e., they need to be well typed in the source. However, λ^{fx} does not have recursive types, so what is the source-level correspondent of $\mu\alpha.\tau$?

We adapt the same intuition of previous work [Devriese et al. 2016, 2017] in our setting too: it is not necessary to precisely embed target types into the source language in order to backtranslate terms. In fact, we need to reason for *up to n steps*, which means that we can approximate target types n -levels deep. Thus, concretely, we do not need recursive types in λ^{fx} . Given a target recursive type, we unfold it n times and backtranslate its unfolding to model the n target reductions required.

According to this strategy, the backtranslation of a term of type τ should have type *unfold τ n times*. During this unfolding, however, things can go wrong. Specifically, we do not know at runtime the level of unfolding we are dealing with, i.e., we cannot inspect n at runtime. Thus, we need a way to model failure (as a sort of catchable exception), or, having reached more than n unfoldings,

because in that case we need to diverge.⁴ Thus at each level of unfolding, we backtranslate τ into “ $\tau \uplus \text{Unit}$ ” (we will make this formal below), where the right Unit models failure.

We make these intuitions concrete and formalise the type for λ_I^μ values backtranslated into λ^{fx} as $\text{BtT}_{n;\tau}^{\text{fl}}$ in Figure 5 (for $\text{Backtranslation Type}$; the superscript indicates the languages involved, the subscripts are effectively parameters of this type). Type $\text{BtT}_{n;\tau}^{\text{fl}}$ is defined inductively on n and it backtranslates the structure of τ in the source type it creates. At no steps ($n=0$), the backtranslation is not needed any more because intuitively we already performed the n steps, so the only type is Unit . Otherwise, the backtranslated type maintains the same structure of the target type. In the case for $\mu\alpha.\tau$, the backtranslated type is the unfolding of $\mu\alpha.\tau$, but at a decremented index (n). Intuitively, this is to match the reduction step that will happen in the target for eliminating $\text{unfold}_{\mu\alpha.\tau} \text{ fold}_{\mu\alpha.\tau}$ annotations.

The type of λ_E^μ terms backtranslated in λ^{fx} ($\text{BtT}_{n;\tau}^{\text{fe}}$, still in Figure 5) has an important difference. The case for $\mu\alpha.\tau$ does not lose a step in the index and simply performs the unfolding of the recursive type without an additional $\uplus \text{Unit}$. This difference matches the fact that in λ_E^μ there is no additional reduction rule in the semantics. Additionally, this difference affects the helper functions needed to deal with values of backtranslation type, as we discuss later.

Intuitively, the fact that the backtranslation of a recursive type is its n -level deep unfolding is possible because $\mu\alpha.\tau$ is contractive in α . This is sufficient because we need to only replicate n steps in order to differentiate terms, so a n -level deep unfolding of the type suffices in order to reach the differentiation. For example, let us take the type of list of booleans in λ_E^μ : $\mu\alpha. \text{Unit} \uplus (\text{Bool} \times \alpha)$ (which we dub List_B) and its first unfolding $\text{Unit} \uplus (\text{Bool} \times \text{List}_B)$ (which we dub List_B^1). The backtranslation (for $n = 3$) for this type is the following: $\text{BtT}_{3;\text{List}_B}^{\text{fe}} = \text{BtT}_{3;\text{Unit} \uplus (\text{Bool} \times \text{List}_B)}^{\text{fe}} = \dots = ((\text{Unit} \uplus \text{Unit}) \uplus (((\text{Bool} \uplus \text{Unit}) \times \text{BtT}_{1;\text{List}_B}^{\text{fe}}) \uplus \text{Unit})) \uplus \text{Unit}$.⁵ Formally, the measure that ensures that this type is well founded is the precision n together with $\text{lmc}(\mu\alpha.\tau)$ i.e., the number of leading μ s in type τ , for reasons analogous to those discussed in Section 2.2.

The type of λ_E^μ terms backtranslated in λ_I^μ ($\text{BtT}_{n;\tau}^{\text{ie}}$) is the same as the one just presented ($\text{BtT}_{n;\tau}^{\text{fe}}$). Intuitively, this is because the n -level deep unfolding of τ in the backtranslation type does not rely on recursive types in λ_I^μ .

3.2.1 Working with the Backtranslation Type. In order to work with values of backtranslated type, we need a way to create and destruct them. Additionally, we need a way to increase and decrease the approximation level (the n index), for reasons we explain below. This is what we present

$$\begin{array}{l}
 \text{BtT}_{0;\tau}^{\text{fl}} \stackrel{\text{def}}{=} \text{Unit} \\
 \text{BtT}_{n+1;\tau}^{\text{fl}} \stackrel{\text{def}}{=} \begin{cases} \text{Unit} \uplus \text{Unit} & \text{if } \tau = \text{Unit} \\ \text{Bool} \uplus \text{Unit} & \text{if } \tau = \text{Bool} \\ (\text{BtT}_{n;\tau}^{\text{fl}} \rightarrow \text{BtT}_{n;\tau'}^{\text{fl}}) \uplus \text{Unit} & \text{if } \tau = \tau \rightarrow \tau' \\ (\text{BtT}_{n;\tau}^{\text{fl}} \times \text{BtT}_{n;\tau'}^{\text{fl}}) \uplus \text{Unit} & \text{if } \tau = \tau \times \tau' \\ (\text{BtT}_{n;\tau}^{\text{fl}} \uplus \text{BtT}_{n;\tau'}^{\text{fl}}) \uplus \text{Unit} & \text{if } \tau = \tau \uplus \tau' \\ \text{BtT}_{n;\tau'}^{\text{fl}}[\mu\alpha.\tau'/\alpha] \uplus \text{Unit} & \text{if } \tau = \mu\alpha.\tau' \end{cases} \\
 \text{BtT}_{0;\tau}^{\text{fe}} \stackrel{\text{def}}{=} \text{Unit} \\
 \text{BtT}_{n+1;\tau}^{\text{fe}} \stackrel{\text{def}}{=} \begin{cases} \text{omitted cases are as above} \\ \text{BtT}_{n+1;\tau'}^{\text{fe}}[\mu\alpha.\tau'/\alpha] \end{cases} \quad \text{if } \tau = \mu\alpha.\tau' \\
 \text{BtT}_{n;\tau}^{\text{ie}} \stackrel{\text{def}}{=} \text{as } \text{BtT}_{n;\tau}^{\text{fe}}
 \end{array}$$

Fig. 5. The type of backtranslated terms (excerpts).

⁴ Recall that one of the two terms ($\llbracket t_1 \rrbracket$ and $\llbracket t_2 \rrbracket$) is guaranteed to terminate within n steps, so if that does not happen, the other term needs to diverge. This ensures that contextually-equivalent terms remain equivalent, i.e., they equi-terminate.

⁵ Where the first $\text{Unit} \uplus \text{Unit}$ is the result of $\text{BtT}_{2;\text{Unit}}^{\text{fe}}$ and the $\text{Bool} \uplus \text{Unit}$ is the result of $\text{BtT}_{1;\text{Bool}}^{\text{fe}}$.

now mainly for terms of type $\text{BtT}_{n;\tau}^{\text{fl}}$, though we report the most interesting cases for the other backtranslation types too. Recall that the definitions of the other two backtranslation types are the same, so these helpers are also the same and we report only one.

Given a target value v of type τ , in order to *create* a source term of type $\text{BtT}_{n;\tau}^{\text{fl}}$ it suffices to create $\text{inl } v$ (informally). However, in order to *use* a source term of type $\text{BtT}_{n;\tau}^{\text{fl}}$ at the expected type τ , we need to destroy it according to τ : this is done by the family of source functions $\text{case}_{n;\tau}^{\text{fl}}$.

$$\text{case}_{n;\tau}^{\text{fl}} = \lambda x : \text{BtT}_{n+1;\tau}^{\text{fl}}. \text{case } x \text{ of } \text{inl } x_1 \mapsto x_1 \mid \text{inr } x_2 \mapsto \text{omega}_{\text{BtT}_{n;\tau}^{\text{fl}}}$$

Intuitively, all these functions strip the value of type $\text{BtT}_{n+1;\tau}^{\text{fl}}$ they take in input of the inl tag and return the underlying value. Thus, at arrow type, the returned value has type $(\text{BtT}_{n;\tau}^{\text{fl}} \rightarrow \text{BtT}_{n;\tau'}^{\text{fl}})$ while at recursive type it has type $\text{BtT}_{n;\tau[\mu\alpha.\tau/\alpha]}^{\text{fl}}$. In case the wrong value is passed in (i.e., it is an inr), these functions diverge via term $\text{omega}_{\text{BtT}_{n;\tau}^{\text{fl}}}$, which is easily encodable in λ^{fx} .

Recall that the $\text{BtT}_{n;\tau}^{\text{fe}}$ for $\tau = \mu\alpha.\tau$ is different: it is just $\text{BtT}_{n;\tau[\mu\alpha.\tau/\alpha]}^{\text{fe}}$ so the type is unfolded and the index is the same. The destructor used for this backtranslation type ($\text{case}_{n;\mu\alpha.\tau}^{\text{fe}}$) is therefore different than the one above. Specifically, we do not need to destruct a backtranslated type indexed with τ because that never arises (i.e., the type is unfolded). Consider type $\text{BtT}_{3;\text{List}_B}^{\text{fe}}$ from before: at index 3 the backtranslation does not handle values of that type but of type $\text{BtT}_{3;\text{List}_B^I}^{\text{fe}}$. That is, it handles values whose top-level connector is the $\uplus$ of List_B . Finally, the destructor used for $\text{BtT}_{n;\mu\alpha.\tau}^{\text{ie}}$ ($\text{case}_{n;\mu\alpha.\tau}^{\text{ie}}$) is analogous to this last one ($\text{case}_{n;\mu\alpha.\tau}^{\text{fe}}$).

$$\text{case}_{n;\tau}^{\text{fe}} = \lambda x : \text{BtT}_{n+1;\tau}^{\text{fe}}. \text{case } x \text{ of } \text{inl } x_1 \mapsto x_1 \mid \text{inr } x_2 \mapsto \text{omega}_{\text{BtT}_{n;\tau}^{\text{fe}}} \quad \tau \neq \mu\alpha.\tau$$

3.2.2 Increasing and Decreasing the Approximation Level. The second piece of formalism that we need is functions to increase or decrease the approximation level of backtranslated terms. We exemplify their necessity with an example from Devriese et al. [2016]. Consider $\lambda x : \tau. \text{inr } x$, intuitively its backtranslation (for a sufficiently-large n) is: $\text{inl } \lambda x : \text{BtT}_{n-1;\tau}^{\text{fl}}. \text{inl } \text{inr } x$. If we try to typecheck this, though, we see that x has type $\text{BtT}_{n-1;\tau}^{\text{fl}}$ while it is expected to have type $\text{BtT}_{n-2;\tau}^{\text{fl}}$, i.e., its index should be lower. This concern is about well-typedness, not precision of the backtranslation. Since x is inside an inr , inspecting it for any number of steps requires at least an additional step, to ‘case’ x out of the inr . In other words, for the inr to be a precise approximation up to $n-1$ steps, x needs to only be precise up to $n-2$ steps. Thus, it is safe to throw away one level of precision and *downgrade* x from type $\text{BtT}_{n-1;\tau}^{\text{fl}}$ to $\text{BtT}_{n-2;\tau}^{\text{fl}}$.

However, downgrading is not sufficient. Consider how we can downgrade a value of type $\text{BtT}_{n+1;\tau \rightarrow \tau'}^{\text{fl}}$ to one of type $\text{BtT}_{n;\tau \rightarrow \tau'}^{\text{fl}}$. We need to convert a function of type $\text{BtT}_{n;\tau}^{\text{fl}} \rightarrow \text{BtT}_{n;\tau'}^{\text{fl}}$ into one of type $\text{BtT}_{n+1;\tau}^{\text{fl}} \rightarrow \text{BtT}_{n+1;\tau'}^{\text{fl}}$. To do this, we need to upgrade the argument value of type $\text{BtT}_{n;\tau}^{\text{fl}}$ into one of type $\text{BtT}_{n+1;\tau}^{\text{fl}}$. Fortunately, this does not mean we need to magically improve the approximation precision of the value concerned. Type $\text{BtT}_{n;\tau}^{\text{fl}}$ has an “error box” ($\cdot \cdot \cdot \uplus \text{Unit}$) at every level so we can simply construct the value such that it simply does not use the additional level of precision in $\text{BtT}_{n;\tau}^{\text{fl}}$.

Finally, another reason we need to upgrade and downgrade a value is that type $\text{BtT}_{n;\tau}^{\text{fl}}$ must be sufficiently large to contain approximations of target values *up to less than n steps*. In fact, for a term to be well-typed the accuracy of the approximation can be less than n . In these cases (i.e., for $m < n$), values of type $\text{BtT}_{n;\tau}^{\text{fl}}$ will be downgraded to type $\text{BtT}_{m;\tau}^{\text{fl}}$. Dually, there will be cases where some values need to be upgraded.

$\text{upgrade}_{n;\tau}^{\text{fl}} : \text{BtT}_{n;\tau}^{\text{fl}} \rightarrow \text{BtT}_{n+1;\tau}^{\text{fl}} \quad \text{and} \quad \text{downgrade}_{n;\tau}^{\text{fl}} : \text{BtT}_{n+1;\tau}^{\text{fl}} \rightarrow \text{BtT}_{n;\tau}^{\text{fl}}$	
$\text{upgrade}_{0;\tau}^{\text{fl}} = \lambda x : \text{BtT}_{0;\tau}^{\text{fl}}. \text{unk}$	
$\text{upgrade}_{n+1;\text{Unit}}^{\text{fl}} = \lambda x : \text{Unit} \uplus \text{Unit}. x$	$\text{unk} = \text{inr unit}$
$\text{upgrade}_{n+1;\tau \times \tau'}^{\text{fl}} = \lambda x : \text{BtT}_{n+1;\tau \times \tau'}^{\text{fl}}.$	
$\text{case } x \text{ of}$	
$\left \begin{array}{l} \text{inl } x_1 \mapsto \text{inl} \left\langle \begin{array}{l} \text{upgrade}_{n;\tau}^{\text{fl}} x_1.1, \\ \text{upgrade}_{n;\tau'}^{\text{fl}} x_1.2 \end{array} \right\rangle \\ \text{inr } x_2 \mapsto \text{inr } x_2 \end{array} \right.$	$\text{downgrade}_{0;\tau}^{\text{fl}} = \lambda x : \text{BtT}_{0;\tau}^{\text{fl}}. \text{unit}$
$\text{upgrade}_{n+1;\tau \rightarrow \tau'}^{\text{fl}} = \lambda x : \text{BtT}_{n+1;\tau \rightarrow \tau'}^{\text{fl}}.$	$\text{downgrade}_{n+1;\text{Unit}}^{\text{fl}} = \lambda x : \text{Unit} \uplus \text{Unit}. x$
$\text{case } x \text{ of}$	$\text{downgrade}_{n+1;\tau \rightarrow \tau'}^{\text{fl}} = \lambda x : \text{BtT}_{n+2;\tau \rightarrow \tau'}^{\text{fl}}.$
$\left \begin{array}{l} \text{inl } x_1 \mapsto \text{inl} \quad \lambda z : \text{BtT}_{n+1;\tau}^{\text{fl}}. \text{upgrade}_{n;\tau'}^{\text{fl}} \\ \quad (x_1 (\text{downgrade}_{n;\tau}^{\text{fl}} z)) \\ \text{inr } x_2 \mapsto \text{inr } x_2 \end{array} \right.$	$\text{case } x \text{ of}$
$\text{upgrade}_{n+1;d\mu\alpha.\tau'}^{\text{fl}} = \lambda x : \text{BtT}_{n+1;d\mu\alpha.\tau'}^{\text{fl}}.$	$\left \begin{array}{l} \text{inl } x_1 \mapsto \text{inl} \quad \lambda z : \text{BtT}_{n;\tau}^{\text{fl}}. \text{downgrade}_{n;\tau'}^{\text{fl}} \\ \quad (x_1 (\text{upgrade}_{n;\tau}^{\text{fl}} z)) \\ \text{inr } x_2 \mapsto \text{inr } x_2 \end{array} \right.$
$\text{case } x \text{ of}$	$\text{downgrade}_{n+1;\mu\alpha.\tau'}^{\text{fl}} = \lambda x : \text{BtT}_{n+2;\mu\alpha.\tau'}^{\text{fl}}.$
$\left \begin{array}{l} \text{inl } x_1 \mapsto \text{inl} (\text{upgrade}_{n;\tau'}^{\text{fl}} [\mu\alpha.\tau' / \alpha] x_1) \\ \text{inr } x_2 \mapsto \text{inr } x_2 \end{array} \right.$	$\text{case } x \text{ of}$
	$\left \begin{array}{l} \text{inl } x_1 \mapsto \text{inl} (\text{downgrade}_{n;\tau'}^{\text{fl}} [\mu\alpha.\tau' / \alpha] x_1) \\ \text{inr } x_2 \mapsto \text{inr } x_2 \end{array} \right.$
$\text{upgrade}_{n+1;\mu\alpha.\tau}^{\text{fe}} = \text{upgrade}_{n+1;\tau[\mu\alpha.\tau / \alpha]}^{\text{fe}}$	$\text{downgrade}_{n+1;\mu\alpha.\tau}^{\text{fe}} = \text{downgrade}_{n+1;\tau[\mu\alpha.\tau / \alpha]}^{\text{fe}}$
$\text{upgrade}_{n;\tau}^{\text{fe}} = \text{as above}$	$\text{downgrade}_{n;\tau}^{\text{fe}} = \text{as above}$
$\text{upgrade}_{n;\tau}^{\text{IE}} = \text{as } \text{upgrade}_{n;\tau}^{\text{fe}}$	$\text{downgrade}_{n;\tau}^{\text{IE}} = \text{as } \text{downgrade}_{n;\tau}^{\text{fe}}$

Fig. 6. Definition of the `upgrade` and `downgrade` functions (excerpts).

Functions `upgradefl` and `downgradefl` perform what we just discussed; their types and formalisation is presented in Figure 6 (partially for space constraints). The cases for `Unit` and `Bool` are optimised based on the fact that $\text{BtT}_{n;\text{Unit}}^{\text{fl}} = \text{BtT}_{m;\text{Unit}}^{\text{fl}}$ (resp. $\text{BtT}_{n;\text{Bool}}^{\text{fl}} = \text{BtT}_{m;\text{Bool}}^{\text{fl}}$) so long as $n, m > 0$. As mentioned, downgrade ‘forgets’ information about the approximation, effectively dropping 1 level of precision in the backtranslation. Dually, upgrade adds 1 level of information in the approximation. Adding this information is, however, not precise, because those additional levels are unknown (`unk`). Effectively, while `downgradefln;τ (upgradefln;τ t)` reduces to `t`, term `upgradefln;τ (downgradefln;τ t)` does not reduce to `t` because information was lost (Example 1).

Example 1 (Upgrading after downgrading forgets information). Consider the following term: `downgradefl0;Bool inl true`, which reduces to `unit`. If we apply `upgradefl0;Bool` to it, we do not obtain back `inl true` but `unk`, which is `inr unit`. That is because downgrade forgets the shape of the value it received (`inl true`) and upgrade cannot possibly recover that information. $\square$

Finally, we need to define these functions for the other backtranslations that rely on the other backtranslation types `BtTfe` and `BtTIE`. As mentioned, the main difference between these last two backtranslation types and `BtTfl` is the case for target recursive types. Recall that these last two

backtranslation types for recursive types perform the unfolding of the type without decrementing the index. This affects these functions too: upgrading or downgrading a term at a recursive type is like upgrading or downgrading at the unfolding of that type but at the same index.

$$\begin{array}{c}
 \hline
 \boxed{\text{in-dn}_{n;\tau}^{\text{fl}} \quad \text{and} \quad \text{case-up}_{n;\tau}^{\text{fl}}} \\
 \hline
 \begin{array}{ll}
 \text{in-dn}_{n;\text{Unit}}^{\text{fl}} = \lambda x : \text{Unit}. \text{downgrade}_{n;\text{Unit}}^{\text{fl}} (\text{inl } x) & \text{in-dn}_{n;\text{Bool}}^{\text{fl}} = \lambda x : \text{Bool}. \text{downgrade}_{n;\text{Bool}}^{\text{fl}} (\text{inl } x) \\
 \text{in-dn}_{n;\tau \rightarrow \tau'}^{\text{fl}} = \lambda x : \text{BtT}_{n;\tau}^{\text{fl}} \rightarrow \text{BtT}_{n;\tau'}^{\text{fl}}. \text{downgrade}_{n;\tau \rightarrow \tau'}^{\text{fl}} (\text{inl } x) & \text{in-dn}_{n;\tau \times \tau'}^{\text{fl}} = \lambda x : \text{BtT}_{n;\tau}^{\text{fl}} \times \text{BtT}_{n;\tau'}^{\text{fl}}. \text{downgrade}_{n;\tau \times \tau'}^{\text{fl}} (\text{inl } x) \\
 \text{in-dn}_{n;\tau \uplus \tau'}^{\text{fl}} = \lambda x : \text{BtT}_{n;\tau}^{\text{fl}} \uplus \text{BtT}_{n;\tau'}^{\text{fl}}. \text{downgrade}_{n;\tau \uplus \tau'}^{\text{fl}} (\text{inl } x) & \text{in-dn}_{n;\mu\alpha.\tau}^{\text{fl}} = \lambda x : \text{BtT}_{n;\tau}^{\text{fl}} [\mu\alpha.\tau / \alpha]. \text{downgrade}_{n;\mu\alpha.\tau}^{\text{fl}} (\text{inl } x) \\
 \text{case-up}_{n;\tau}^{\text{fl}} = \lambda x : \text{BtT}_{n;\tau}^{\text{fl}}. \text{case}_{n;\tau}^{\text{fl}} (\text{upgrade}_{n;\tau}^{\text{fl}} (x))
 \end{array} \\
 \hline
 \text{in-dn}_{n;\tau}^{\text{fe}} \quad \text{and} \quad \text{case-up}_{n;\tau}^{\text{fe}} = \text{as above, without a case for } \tau = \mu\alpha.\tau \\
 \hline
 \text{in-dn}_{n;\tau}^{\text{IE}} \quad \text{and} \quad \text{case-up}_{n;\tau}^{\text{IE}} = \text{as above, without a case for } \tau = \mu\alpha.\tau \\
 \hline
 \end{array}$$

Fig. 7. Compacted functions used to manipulate backtranslated values.

In the backtranslation, we generally use creation of a backtranslated value together with a $\text{downgrade}^{\text{fl}}$, while we use destruction of backtranslated values together with an $\text{upgrade}^{\text{fl}}$. Thus, we provide compacted functions that do exactly this, $\text{in-dn}_{n;\tau}^{\text{fl}}$ and $\text{case-up}_{n;\tau}^{\text{fl}}$ (Figure 7). Note that the arguments to the first function is not ill-typeset: they indeed take a parameter whose type is the inl projection of type $\text{BtT}_{n;\text{Unit}}^{\text{fl}}$. As for the previous helpers, the compacted versions that operate on terms of type $\text{BtT}_{n;\mu\alpha.\tau}^{\text{fe}}$ (and $\text{BtT}_{n;\mu\alpha.\tau}^{\text{IE}}$) are different. Since there is no destructor for $\text{BtT}_{n;\mu\alpha.\tau}^{\text{fe}}$, there also is no need for a compacted version.

At this point we may ask ourselves: how can we reason about these functions, as well as about backtranslated terms? This is what we explain next.

3.3 Relating Backtranslated Terms

If we were to use the logical relations of Figure 3 to relate a term and its backtranslation, this would simply not work. Consider λ_1^μ type **Unit**, that is backtranslated (at any approximation $n > 0$) into $\text{BtT}_{n;\text{Unit}}^{\text{fl}}$, i.e., $\text{Unit} \uplus \text{Unit}$. Value **unit** should normally be backtranslated to inl unit . Following the value relation in LR_{μ}^{fx} for $\uplus$ types, both terms need to have an inl tag, so this does not work. More importantly, it *should not* work: we are not relating terms of $\uplus$ type, we are relating backtranslated terms, where the backtranslation performs a modification on the type (and thus the term) by inserting the inl .

This is the reason we have pseudotypes and, in particular, the reason we have EmulT . We have three EmulT s—one per backtranslation—and each follows the same intuition, which we explain starting with $\text{EmulT}_{n;p;\tau}^{\text{fl}}$, the type of backtranslated λ_1^μ terms into λ^{fx} (top of Figure 8). $\text{EmulT}_{n;p;\tau}^{\text{fl}}$ is indexed by a non-negative number n , a value $p ::= \text{precise} \mid \text{imprecise}$ and the original target type τ . The number tracks the depth of type that are being related, index p tracks the precision of the approximation (as explained below) and the original type carries precise information of the type to expect in the backtranslation. As seen, sometimes we have **unk** values (i.e., **inr unit**)

$$\begin{array}{c}
\mathcal{V} \left[\text{EmulT}_{0;\text{imprecise};\tau}^{\text{fl}} \right]_{\nabla} \stackrel{\text{def}}{=} \{(W, v, v) \mid v = \text{unit}\} \qquad \mathcal{V} \left[\text{EmulT}_{0;\text{precise};\tau}^{\text{fl}} \right]_{\nabla} \stackrel{\text{def}}{=} \emptyset \\
\mathcal{V} \left[\text{EmulT}_{n+1;p;\tau}^{\text{fl}} \right]_{\nabla} \stackrel{\text{def}}{=} \{(W, v, v) \mid v \in \text{oftype}(\text{EmulT}_{n+1;p;\tau}^{\text{fl}}) \text{ and } v \in \text{oftype}(\tau) \text{ and} \\
\text{either } \cdot v = \text{inr unit} \text{ and } p = \text{imprecise} \\
\text{or } \cdot \left. \begin{array}{l}
\tau = \text{Unit} \text{ and } \exists v'. v = \text{inl } v' \text{ and } (W, v', v) \in \mathcal{V} \left[\text{Unit} \right]_{\nabla} \\
\tau = \text{Bool} \text{ and } \exists v'. v = \text{inl } v' \text{ and } (W, v', v) \in \mathcal{V} \left[\text{Bool} \right]_{\nabla} \\
\tau = \tau_1 \rightarrow \tau_2 \text{ and } \exists v'. v = \text{inl } v' \text{ and } (W, v', v) \in \mathcal{V} \left[\text{EmulT}_{n;p;\tau_1}^{\text{fl}} \rightarrow \text{EmulT}_{n;p;\tau_2}^{\text{fl}} \right]_{\nabla} \\
\tau = \tau_1 \times \tau_2 \text{ and } \exists v'. v = \text{inl } v' \text{ and } (W, v', v) \in \mathcal{V} \left[\text{EmulT}_{n;p;\tau_1}^{\text{fl}} \times \text{EmulT}_{n;p;\tau_2}^{\text{fl}} \right]_{\nabla} \\
\tau = \tau_1 \uplus \tau_2 \text{ and } \exists v'. v = \text{inl } v' \text{ and } (W, v', v) \in \mathcal{V} \left[\text{EmulT}_{n;p;\tau_1}^{\text{fl}} \uplus \text{EmulT}_{n;p;\tau_2}^{\text{fl}} \right]_{\nabla} \\
\tau = \mu\alpha. \tau \text{ and } \exists v'. v = \text{inl } v' \text{ and} \\
\exists v'. v = \text{fold}_{\mu\alpha. \tau} v' (W, v', v') \in \triangleright \mathcal{V} \left[\text{EmulT}_{n;p;\tau[\mu\alpha. \tau/\alpha]}^{\text{fl}} \right]_{\nabla}
\end{array} \right\} \\
\hline
\mathcal{V} \left[\text{EmulT}_{0;\text{imprecise};\tau}^{\text{fe}} \right]_{\nabla} \stackrel{\text{def}}{=} \{(W, v, v) \mid v = \text{unit}\} \qquad \mathcal{V} \left[\text{EmulT}_{0;\text{precise};\tau}^{\text{fe}} \right]_{\nabla} \stackrel{\text{def}}{=} \emptyset \\
\mathcal{V} \left[\text{EmulT}_{n+1;p;\tau}^{\text{fe}} \right]_{\nabla} \stackrel{\text{def}}{=} \{(W, v, v) \mid v \in \text{oftype}(\text{EmulT}_{n+1;p;\tau}^{\text{fe}}) \text{ and } v \in \text{oftype}(\tau) \text{ and} \\
\text{either } \cdot v = \text{inr unit} \text{ and } p = \text{imprecise} \\
\text{or } \cdot \left. \begin{array}{l}
\text{omitted parts are as above} \\
\tau = \mu\alpha. \tau \text{ and } \tau \text{ contractive in } \alpha \text{ and } (W, v, v) \in \mathcal{V} \left[\text{EmulT}_{n+1;p;\tau[\mu\alpha. \tau/\alpha]}^{\text{fl}} \right]_{\nabla}
\end{array} \right\} \\
\hline
\mathcal{V} \left[\text{EmulT}_{n;p;\tau}^{\text{IE}} \right]_{\nabla} \text{ is defined analogously to } \mathcal{V} \left[\text{EmulT}_{n;p;\tau}^{\text{fe}} \right]_{\nabla} \\
\hline
\begin{array}{ll}
\text{repEmul}^{\text{fI}}(\text{EmulT}_{n;p;\tau}^{\text{fl}}) = \text{BtT}_{n;\tau}^{\text{fl}} & \text{repEmul}^{\text{fI}}(\hat{r}_1 \rightarrow \hat{r}_2) = \text{repEmul}^{\text{fI}}(\hat{r}_1) \rightarrow \text{repEmul}^{\text{fI}}(\hat{r}_2) \\
\text{repEmul}^{\text{fI}}(\text{Bool}) = \text{Bool} & \text{repEmul}^{\text{fI}}(\hat{r}_1 \times \hat{r}_2) = \text{repEmul}^{\text{fI}}(\hat{r}_1) \times \text{repEmul}^{\text{fI}}(\hat{r}_2) \\
\text{repEmul}^{\text{fI}}(\text{Unit}) = \text{Unit} & \text{repEmul}^{\text{fI}}(\hat{r}_1 \uplus \hat{r}_2) = \text{repEmul}^{\text{fI}}(\hat{r}_1) \uplus \text{repEmul}^{\text{fI}}(\hat{r}_2) \\
\text{fxToIs}(\text{EmulT}_{n;p;\tau}^{\text{fl}}) = \tau & \text{fxToIs}(\hat{r}_1 \rightarrow \hat{r}_2) = \text{fxToIs}(\hat{r}_1) \rightarrow \text{fxToIs}(\hat{r}_2) \\
\text{fxToIs}(\text{Unit}) = \text{Unit} & \text{fxToIs}(\hat{r}_1 \times \hat{r}_2) = \text{fxToIs}(\hat{r}_1) \times \text{fxToIs}(\hat{r}_2) \\
\text{fxToIs}(\text{Bool}) = \text{Bool} & \text{fxToIs}(\hat{r}_1 \uplus \hat{r}_2) = \text{fxToIs}(\hat{r}_1) \uplus \text{fxToIs}(\hat{r}_2) \\
\text{repEmul}^{\text{fE}}(\cdot) : \hat{t} \rightarrow \tau & \text{repEmul}^{\text{IE}}(\cdot) : \hat{t} \rightarrow \tau \quad \text{fxToEq}(\cdot) : \hat{t} \rightarrow \tau \quad \text{isToEq}(\cdot) : \hat{t} \rightarrow \tau
\end{array}
\end{array}$$

Fig. 8. Missing bits of the logical relation: value relation for backtranslation type (excerpts). Note that p can be either **precise** or **imprecise** in the second clause (the 'or') of the $n+1$ case.

in the backtranslation. Thus, $\mathcal{V} \left[\text{EmulT}_{n;p;\tau}^{\text{fl}} \right]_{\nabla}$ regulates how these values occur depending on the precision index. $p = \text{imprecise}$ will only be used in the $\lesssim$ direction of the approximation, i.e., we have that source termination in *any* number of steps implies target termination. Here, $\mathcal{V} \left[\text{EmulT}_{n;p;\tau}^{\text{fl}} \right]_{\nabla}$ allows **unk** values to occur anywhere in a backtranslated term, and they can correspond to arbitrary target terms. These constraints are simple to enforce because with $\lesssim$ we can achieve this by making backtranslated terms diverge whenever they try to use a **unk** value. This is sufficient because the $\lesssim$ approximation trivially holds when the source term diverges.

On the other hand, $p = \text{precise}$ will be used for the other direction of approximation: $\gtrsim$. Recall that for this direction, termination of target terms in less than n steps implies termination of source

terms. In this case, the requirements on backtranslated terms are stronger: **unk** is ruled out by the definition of $\mathcal{V} \llbracket \text{EmulT}_{n;p;\tau}^{\text{fl}} \rrbracket_{\nabla}$ within depth n , i.e., we cannot reach **unk** in the steps of the world.

The pseudotype for the λ_E^μ to λ^{fx} backtranslation (EmulT^{fE}) follows the same pattern as BtT^{fE} : it does not lose a step in the $\mu\alpha.\tau$ case (Figure 8). At a cursory glance, it appears that a non-contractive $\mu\alpha.\tau$ ruins the well-foundedness of our induction as without decrementing our step index, a non-contractive type seems to infinitely recurse under this definition. Fortunately, however, the condition $v \in \text{oftype}(\tau)$, which with the fact that no values exist of non-contractive types prevents this concern from arising. As before, the pseudotype for the λ_E^μ to λ_1^μ backtranslation (EmulT^{fE}) follows the same approach as EmulT^{fE} .

Finally, we can define function $\text{repEmul}^{\text{fI}}(\cdot)$ that translate from source pseudo-types into plain source types and function $\text{fxToIs}(\cdot)$, that translates source pseudotypes into target types. We present the formalisation for the case for source types being λ^{fx} types and target types being λ_1^μ types. As expected, these functions exists for all backtranslations and they follow the same pattern presented here; for completeness we only report the names and types of the omitted ones.

4 THE THREE COMPILERS AND THEIR BACKTRANSLATIONS

Our compilers (Section 4.1) and backtranslations (Section 4.2) translate between languages as depicted in Figure 1. After showing their formalisation and proving that they relate terms cross-language, this section proves the compilers are fully abstract (Section 4.3).

4.1 Compilers and Reflection of Fully-Abstract Compilation

The compilers (Figure 9) are all mostly homomorphic apart from what we describe below. We overload the compilation notation and express the compiler for types and terms in the same way (we omit the compiler for types since it is the identity). Compiler $\llbracket \cdot \rrbracket_{\lambda_1^\mu}^{\lambda^{\text{fx}}}$ translates **fix** into the Z-combinator annotated with **fold** and **unfold** for λ_1^μ . We cannot use the Y combinator since it does not work in call-by-value [Devriese et al. 2017; New et al. 2016], but fortunately the Z-combinator does [Pierce 2002, Sec. 5]. Compiler $\llbracket \cdot \rrbracket_{\lambda_E^\mu}^{\lambda_1^\mu}$ erases **fold** and **unfold** annotations since λ_E^μ does not have them. Compiler $\llbracket \cdot \rrbracket_{\lambda_E^\mu}^{\lambda^{\text{fx}}}$ is just the composition of the previous two.

Correctness of the compilation (Lemmas 4 to 6 below) is proven via a series of standard compatibility lemmas (Lemma 3, we report just the case for lambda since the others follow the same structure). These, in turn, rely on a series of standard results for these kinds of logical relations such as the fact that related terms plugged in related contexts are still related and antireduction (i.e., if two terms step to related terms, then they are themselves related).

LEMMA 3 (COMPATIBILITY FOR λ). *if $\Gamma, x : \tau' \vdash t \nabla_n t : \tau$ then $\Gamma \vdash \lambda x : \tau'. t \nabla_n \lambda x : \tau'. t : \tau' \rightarrow \tau$*

LEMMA 4 ($\llbracket \cdot \rrbracket_{\lambda_1^\mu}^{\lambda^{\text{fx}}}$ IS SEMANTICS PRESERVING). *if $\Gamma \vdash t : \tau$ then $\Gamma \vdash t \nabla_n \llbracket t \rrbracket_{\lambda_1^\mu}^{\lambda^{\text{fx}}} : \tau$*

LEMMA 5 ($\llbracket \cdot \rrbracket_{\lambda_E^\mu}^{\lambda^{\text{fx}}}$ IS SEMANTICS PRESERVING). *if $\Gamma \vdash t : \tau$ then $\Gamma \vdash t \nabla_n \llbracket t \rrbracket_{\lambda_E^\mu}^{\lambda^{\text{fx}}} : \tau$*

LEMMA 6 ($\llbracket \cdot \rrbracket_{\lambda_E^\mu}^{\lambda_1^\mu}$ IS SEMANTICS PRESERVING). *if $\Gamma \vdash t : \tau$ then $\Gamma \vdash t \nabla_n \llbracket t \rrbracket_{\lambda_E^\mu}^{\lambda_1^\mu} : \tau$*

Since fully-abstract compilation requires reasoning about program contexts, we extend the compiler to operate on them too. This follows the same structure of the compilers above and therefore we omit this definition. Correctness of the compiler scales to contexts too (Lemma 7).

Fig. 9. Definition of our compilers (excerpts).

LEMMA 7 ($[\![\cdot]\!]_{\lambda_1^\mu}^{\lambda^{fx}}$ IS SEMANTICS PRESERVING FOR CONTEXTS).

if $\vdash \mathcal{C} : \Gamma, \tau \rightarrow \Gamma', \tau'$ then $\vdash \mathcal{C} \nabla_n [\![\mathcal{C}]\!]_{\lambda_1^\mu}^{\lambda^{fx}} : \Gamma, \tau \rightarrow \Gamma', \tau'$

With these results, we can already prove the reflection direction of fully-abstract compilation (Theorems 8 to 10). The proof follows the structure depicted in the left part of Figure 4.

Theorem 8 ($[\![\cdot]\!]_{\lambda_1^\mu}^{\lambda^{fx}}$ reflects equivalence). If $\emptyset \vdash [\![t_1]\!]_{\lambda_1^\mu}^{\lambda^{fx}} \simeq_{\text{ctx}} [\![t_2]\!]_{\lambda_1^\mu}^{\lambda^{fx}} : [\![\tau]\!]_{\lambda_1^\mu}^{\lambda^{fx}}$ then $\emptyset \vdash t_1 \simeq_{\text{ctx}} t_2 : \tau$

Theorem 9 ($[\![\cdot]\!]_{\lambda_E^\mu}^{\lambda_1^\mu}$ reflects equivalence). If $\emptyset \vdash [\![t_1]\!]_{\lambda_E^\mu}^{\lambda_1^\mu} \simeq_{\text{ctx}} [\![t_2]\!]_{\lambda_E^\mu}^{\lambda_1^\mu} : [\![\tau]\!]_{\lambda_E^\mu}^{\lambda_1^\mu}$ then $\emptyset \vdash t_1 \simeq_{\text{ctx}} t_2 : \tau$

Theorem 10 ($[\![\cdot]\!]_{\lambda_E^\mu}^{\lambda^{fx}}$ reflects equivalence). If $\emptyset \vdash [\![t_1]\!]_{\lambda_E^\mu}^{\lambda^{fx}} \simeq_{\text{ctx}} [\![t_2]\!]_{\lambda_E^\mu}^{\lambda^{fx}} : [\![\tau]\!]_{\lambda_E^\mu}^{\lambda^{fx}}$ then $\emptyset \vdash t_1 \simeq_{\text{ctx}} t_2 : \tau$

Since this last compiler is the composition of the other two, the proof of Theorem 10 trivially follows from composing the proofs of the other two compilers.

4.2 Backtranslations and Preservation of Fully-Abstract Compilation

Function $\text{emulate}^{\text{fl}}(\cdot)$ is responsible for translating a target term of type τ into a source one of type $\text{BtT}_{n;\tau}^{\text{fl}}$ (Section 4.2.1) by relying on the machinery needed for working with BtT^{fl} terms from Section 3.2. This function is easily extended to work with program contexts, producing contexts with hole of type $\text{BtT}_{n;\tau}^{\text{fl}}$. However, recall that the goal of the backtranslation is generating a source context whose hole can be filled with source terms t_1 and t_2 and their type is not $\text{BtT}_{n;\tau}^{\text{fl}}$ but τ . Thus, there is a mismatch between the type of the hole of the emulated context and that of the terms to be plugged there. Since emulated contexts work with BtT^{fl} values, we need a function that wraps terms of an arbitrary type τ into a value of type $\text{BtT}_{n;\tau}^{\text{fl}}$. This function is called $\text{inject}^{\text{fl}}$ (Section 4.2.2) and it is the last addition we need before the backtranslations (Section 4.2.3).

4.2.1 Emulation of Terms and Contexts. Like the compiler, the emulation must not just operate on types and terms, but also on program contexts. Unlike the compiler, the emulation operates on *type derivations* for terms and contexts since all our target languages are typed. Thus, the emulation of a lambda would look like the following (using D as a metavariable to range over derivations and omitting functions to work with BtT^{fl}).

$$\text{emulate}^{\text{fl}}\left(\frac{\frac{D}{\Gamma, x : \tau \vdash t : \tau'}}{\Gamma \vdash \lambda x : \tau. t : \tau \rightarrow \tau'}\right) = \lambda x : \text{BtT}_{n;\tau}^{\text{fl}}. \text{emulate}^{\text{fl}}\left(\frac{D}{\Gamma, x : \tau \vdash t : \tau'}\right)$$

However, note that each judgement uniquely identifies which typing rule is being applied and the underlying derivation. Thus, for compactness, we only write the judgement in the emulation and implicitly apply the related typing rule to obtain the underlying judgements for recursive calls.

Function $\text{emulate}_n^{\text{fl}}(\cdot)$ (Figures 10 and 11) is indexed by the approximation index n in order to know which BtT^{fl} -helper functions to use. There are few interesting bits in the emulation of terms (and of contexts). When emulating constructors for terms of type τ , we create a value of the corresponding backtranslation type $\text{BtT}_{n;\tau}^{\text{fl}}$ and, in order to be well-typed, we $\text{downgrade}^{\text{fl}}$ that value by 1. Dually, emulating destructors for terms of type τ requires upgrading the term for 1 level of precision because they are then destructed to access the underlying type. When emulating λ_1^μ derivations into λ^{fx} , we need to consider the case when $\text{fold}_{\mu\alpha.\tau}$ and $\text{unfold}_{\mu\alpha.\tau}$ annotations are encountered. There, we know that the backtranslation will work with terms typed at the unfolding of $\mu\alpha.\tau$, so we simply perform the recursive call and insert the appropriate helper function to ensure the resulting term is well-typed.

When emulating λ_E^μ derivations (in the other two emulates in Figure 10), we need to consider the case when term t is given type τ knowing it had type σ and that $\sigma \doteq \tau$ (Rule $\lambda_E^\mu\text{-Type-eq}$). Here we rely on a crucial observation: given two equivalent types, their backtranslation types are *the same* (Theorem 11). To understand why this is the case, consider how the definition of $\text{BtT}_{n;\tau}^{\text{fl}}$ simply unfolds recursive types without losing precision, i.e. it essentially only looks at the depth- n unfolding of type τ and these unfoldings are equal for equal types $\tau \doteq \sigma$. With this fact, we can get away with just performing the recursive call on the sub-derivation for t at type σ .

Theorem 11 (Equivalent types are backtranslated to the same type). If $\tau \doteq \sigma$ then $\text{BtT}_{n;\tau}^{\text{fe}} = \text{BtT}_{n;\sigma}^{\text{fe}}$

Finally, consider $\text{emulate}^{\text{IE}}(\cdot)$, i.e., the emulation of λ_E^μ terms into λ_1^μ : there is no construct that adds fold/unfold annotations. This is due to the same intuition presented before regarding the unfolding of the backtranslation type $\text{BtT}_{n;\mu\alpha.\tau}^{\text{IE}}$, which is $\text{BtT}_{n;\tau}^{\text{IE}}[\mu\alpha.\tau/\alpha]$ i.e. the indexing type is unfolded but the step is not decreased. Intuitively, the backtranslation performs an n -level deep unfolding of the recursive types and operates on those. Thus, backtranslated contexts do not use recursive types but just their n -level deep unfolding, so their annotations are not needed.

$\text{emulate}_n^{\text{fl}}(\cdot) : \Gamma \vdash t : \tau \rightarrow t$

$$\begin{aligned}
\text{emulate}_n^{\text{fl}}(\Gamma \vdash \text{unit} : \text{Unit}) &\stackrel{\text{def.}}{=} \text{in-dn}_{n;\text{Unit}}^{\text{fl}} \text{unit} & \text{emulate}_n^{\text{fl}}(\Gamma \vdash \text{true} : \text{Bool}) &\stackrel{\text{def.}}{=} \text{in-dn}_{n;\text{Bool}}^{\text{fl}} \text{true} \\
\text{emulate}_n^{\text{fl}}(\Gamma \vdash \text{false} : \text{Bool}) &\stackrel{\text{def.}}{=} \text{in-dn}_{n;\text{Bool}}^{\text{fl}} \text{false} & \text{emulate}_n^{\text{fl}}(\Gamma \vdash x : \tau) &\stackrel{\text{def.}}{=} x \\
\text{emulate}_n^{\text{fl}}(\Gamma \vdash \lambda x : \tau. t : \tau \rightarrow \tau') &\stackrel{\text{def.}}{=} \text{in-dn}_{n;\tau \rightarrow \tau'}^{\text{fl}} (\lambda x : \text{BtT}_{n;\tau}^{\text{fl}}. \text{emulate}_n^{\text{fl}}(\Gamma, x : \tau \vdash t : \tau')) \\
\text{emulate}_n^{\text{fl}}(\Gamma \vdash t : t' : \tau) &\stackrel{\text{def.}}{=} (\text{case-up}_{n;\tau' \rightarrow \tau}^{\text{fl}} \text{emulate}_n^{\text{fl}}(\Gamma \vdash t : \tau' \rightarrow \tau)) (\text{emulate}_n^{\text{fl}}(\Gamma \vdash t' : \tau')) \\
\text{emulate}_n^{\text{fl}}(\Gamma \vdash \langle t, t' \rangle : \tau \times \tau') &\stackrel{\text{def.}}{=} \text{in-dn}_{n;\tau \times \tau'}^{\text{fl}} \langle \text{emulate}_n^{\text{fl}}(\Gamma \vdash t : \tau), \text{emulate}_n^{\text{fl}}(\Gamma \vdash t' : \tau') \rangle \\
\text{emulate}_n^{\text{fl}}(\Gamma \vdash t.1 : \tau) &\stackrel{\text{def.}}{=} (\text{case-up}_{n;\tau \times \tau'}^{\text{fl}} \text{emulate}_n^{\text{fl}}(\Gamma \vdash t : \tau \times \tau')).1 \\
\text{emulate}_n^{\text{fl}}(\Gamma \vdash t.2 : \tau) &\stackrel{\text{def.}}{=} (\text{case-up}_{n;\tau' \times \tau}^{\text{fl}} \text{emulate}_n^{\text{fl}}(\Gamma \vdash t : \tau' \times \tau)).2 \\
\text{emulate}_n^{\text{fl}}\left(\Gamma \vdash \text{case } t \text{ of } \begin{array}{l} \text{inl } x_1 \mapsto t' \\ \text{inr } x_2 \mapsto t'' : \tau \end{array}\right) &\stackrel{\text{def.}}{=} \text{case } \left(\text{case-up}_{n;\tau_1 \uplus \tau_2}^{\text{fl}} \text{emulate}_n^{\text{fl}}(\Gamma \vdash t : \tau_1 \uplus \tau_2) \right) \\
&\quad \text{of } \begin{array}{l} \text{inl } x_1 \mapsto \text{emulate}_n^{\text{fl}}(\Gamma, (x_1 : \tau_1) \vdash t' : \tau) \\ \text{inr } x_2 \mapsto \text{emulate}_n^{\text{fl}}(\Gamma, (x_2 : \tau_2) \vdash t'' : \tau) \end{array} \\
\text{emulate}_n^{\text{fl}}(\Gamma \vdash \text{inl } t : \tau \uplus \tau') &\stackrel{\text{def.}}{=} \text{in-dn}_{n;\tau \uplus \tau'}^{\text{fl}} (\text{inl } \text{emulate}_n^{\text{fl}}(\Gamma \vdash t : \tau)) \\
\text{emulate}_n^{\text{fl}}(\Gamma \vdash \text{inr } t : \tau \uplus \tau') &\stackrel{\text{def.}}{=} \text{in-dn}_{n;\tau \uplus \tau'}^{\text{fl}} (\text{inr } \text{emulate}_n^{\text{fl}}(\Gamma \vdash t : \tau')) \\
\text{emulate}_n^{\text{fl}}\left(\Gamma \vdash \begin{array}{l} \text{if } t \text{ then } t_1 \\ \text{else } t_2 : \tau \end{array}\right) &\stackrel{\text{def.}}{=} \text{if } (\text{case-up}_{n;\text{Bool}}^{\text{fl}} \text{emulate}_n^{\text{fl}}(\Gamma \vdash t : \text{Bool})) \\
&\quad \text{then } \text{emulate}_n^{\text{fl}}(\Gamma \vdash t_1 : \tau) \text{ else } \text{emulate}_n^{\text{fl}}(\Gamma \vdash t_2 : \tau) \\
\text{emulate}_n^{\text{fl}}(\Gamma \vdash t : t' : \tau) &\stackrel{\text{def.}}{=} (\text{case-up}_{n;\text{Unit}}^{\text{fl}} \text{emulate}_n^{\text{fl}}(\Gamma \vdash t : \text{Unit})) ; \text{emulate}_n^{\text{fl}}(\Gamma \vdash t' : \tau) \\
\text{emulate}_n^{\text{fl}}(\Gamma \vdash \text{fold}_{\mu\alpha.\tau} t : \mu\alpha.\tau) &\stackrel{\text{def.}}{=} \text{in-dn}_{n;\tau[\mu\alpha.\tau/\alpha]}^{\text{fl}} \text{emulate}_n^{\text{fl}}(\Gamma \vdash t : \tau[\mu\alpha.\tau/\alpha]) \\
\text{emulate}_n^{\text{fl}}\left(\Gamma \vdash \begin{array}{l} \text{unfold}_{\mu\alpha.\tau} t \\ : \tau[\mu\alpha.\tau/\alpha] \end{array}\right) &\stackrel{\text{def.}}{=} \text{case-up}_{n;\mu\alpha.\tau}^{\text{fl}} \text{emulate}_n^{\text{fl}}(\Gamma \vdash t : \mu\alpha.\tau)
\end{aligned}$$

$$\text{emulate}_n^{\text{fE}}\left(\frac{\Gamma \vdash t : \tau \quad \tau \stackrel{\circ}{=} \sigma}{\Gamma \vdash t : \sigma}\right) \stackrel{\text{def.}}{=} \text{emulate}_n^{\text{fE}}(\Gamma \vdash t : \tau) \quad \text{emulate}_n^{\text{fE}}(\dots) \stackrel{\text{def.}}{=} \begin{array}{l} \text{other cases} \\ \text{are as above} \end{array}$$

$$\text{emulate}_n^{\text{IE}}(\dots) \stackrel{\text{def.}}{=} \text{as } \text{emulate}_n^{\text{fE}}(\dots)$$

Fig. 10. Emulation of target terms into source ones.

In order to state that $\text{emulate}^{\text{fl}}(\cdot)$ is correct, we rely on compatibility lemmas akin to those used for compiler correctness (recall Lemma 3). First, note that all our logical relations relate a source and target term at a source pseudo-type. We have extended the logical relation to express the relation between a source and target term at pseudotype EmulT^{fl} , so we should use this to relate a target term and its backtranslation. Second, all logical relations require a source environment to relate terms, and in this case we are given a target environment (the one for the typing of the backtranslated term). To create a source environment starting from this target environment, we take each bound variable and give it backtranslation type using function $\text{toEmul}(\cdot)$. Finally, in these lemmas we need to account for the different directions of the approximation we have. Thus, these compatibility lemmas require that either $n < m$ (so that the results only hold in worlds W with

$$\text{emulate}_n^{\text{fl}}(\cdot) : (\vdash \mathcal{C} : \Gamma, \tau \rightarrow \Gamma', \tau') \rightarrow \mathcal{C}$$

$$\begin{aligned} \text{emulate}_n^{\text{fl}}([\cdot]) &\stackrel{\text{def}}{=} [\cdot] \\ \text{emulate}_n^{\text{fl}}\left(\vdash \lambda x : \tau', \mathcal{C} : \Gamma'', \tau'' \rightarrow \Gamma, \tau' \rightarrow \tau\right) &\stackrel{\text{def}}{=} \text{in-dn}_{n;\tau \rightarrow \tau'}^{\text{fl}}\left(\lambda x : \text{BtT}_{n;\tau}^{\text{fl}}. \text{emulate}_n^{\text{fl}}(\vdash \mathcal{C} : \Gamma'', \tau'' \rightarrow \Gamma, x : \tau', \tau)\right) \\ \text{emulate}_n^{\text{fl}}(\vdash \mathcal{C} \ t_2 : \Gamma', \tau' \rightarrow \Gamma, \tau_2) &\stackrel{\text{def}}{=} \left(\text{case-up}_{n;\tau' \rightarrow \tau}^{\text{fl}} \text{emulate}_n^{\text{fl}}(\vdash \mathcal{C} : \Gamma', \tau' \rightarrow \Gamma, \tau_1 \rightarrow \tau_2)\right) \\ &\quad \left(\text{emulate}_n^{\text{fl}}(\Gamma \vdash t_2 : \tau_1)\right) \\ \text{emulate}_n^{\text{fl}}(\vdash t_1 \ \mathcal{C} : \Gamma', \tau' \rightarrow \Gamma, \tau_2) &\stackrel{\text{def}}{=} \left(\text{case-up}_{n;\tau \rightarrow \tau'}^{\text{fl}} \text{emulate}_n^{\text{fl}}(\Gamma \vdash t_1 : \tau_1 \rightarrow \tau_2)\right) \\ &\quad \left(\text{emulate}_n^{\text{fl}}(\vdash \mathcal{C} : \Gamma', \tau' \rightarrow \Gamma, \tau_1)\right) \\ \text{emulate}_n^{\text{fl}}(\vdash \mathcal{C}.1 : \Gamma', \tau' \rightarrow \Gamma, \tau_1) &\stackrel{\text{def}}{=} \left(\text{case-up}_{n;\tau \times \tau'}^{\text{fl}} \text{emulate}_n^{\text{fl}}(\vdash \mathcal{C} : \Gamma', \tau' \rightarrow \Gamma, \tau_1 \times \tau_2)\right).2 \\ \text{emulate}_n^{\text{fl}}(\vdash \mathcal{C}.2 : \Gamma', \tau' \rightarrow \Gamma, \tau_2) &\stackrel{\text{def}}{=} \left(\text{case-up}_{n;\tau \times \tau'}^{\text{fl}} \text{emulate}_n^{\text{fl}}(\vdash \mathcal{C} : \Gamma', \tau' \rightarrow \Gamma, \tau_1 \times \tau_2)\right).1 \\ \text{emulate}_n^{\text{fl}}\left(\vdash \langle \mathcal{C}, t_2 \rangle : \Gamma', \tau' \rightarrow \Gamma, \tau_1 \times \tau_2\right) &\stackrel{\text{def}}{=} \text{in-dn}_{n;\tau_1 \times \tau_2}^{\text{fl}}\left(\langle \text{emulate}_n^{\text{fl}}(\vdash \mathcal{C} : \Gamma', \tau' \rightarrow \Gamma, \tau_1), \text{emulate}_n^{\text{fl}}(\Gamma \vdash t_2 : \tau_2) \rangle\right) \\ \text{emulate}_n^{\text{fl}}\left(\vdash \langle t_1, \mathcal{C} \rangle : \Gamma', \tau' \rightarrow \Gamma, \tau_1 \times \tau_2\right) &\stackrel{\text{def}}{=} \text{in-dn}_{n;\tau_1 \times \tau_2}^{\text{fl}}\left(\langle \text{emulate}_n^{\text{fl}}(\Gamma \vdash t_1 : \tau_1), \text{emulate}_n^{\text{fl}}(\vdash \mathcal{C} : \Gamma', \tau' \rightarrow \Gamma, \tau_2) \rangle\right) \\ \text{emulate}_n^{\text{fl}}(\vdash \text{fold}_{\mu\alpha.\tau} \mathcal{C} : \Gamma', \tau' \rightarrow \Gamma, \mu\alpha.\tau) &\stackrel{\text{def}}{=} \text{in-dn}_{n;\tau[\mu\alpha.\tau/\alpha]}^{\text{fl}} \text{emulate}_n^{\text{fl}}(\vdash \mathcal{C} : \Gamma', \tau' \rightarrow \Gamma, \tau[\mu\alpha.\tau/\alpha]) \\ \text{emulate}_n^{\text{fl}}(\vdash \text{unfold}_{\mu\alpha.\tau} \mathcal{C} : \Gamma', \tau' \rightarrow \Gamma, \tau[\mu\alpha.\tau/\alpha]) &\stackrel{\text{def}}{=} \text{case-up}_{n;\mu\alpha.\tau}^{\text{fl}} \text{emulate}_n^{\text{fl}}(\vdash \mathcal{C} : \Gamma', \tau' \rightarrow \Gamma, \mu\alpha.\tau) \end{aligned}$$

Fig. 11. Emulation of target contexts into source ones (excerpts).

$\text{lev}(W) \leq n < m$) and $p = \text{precise}$ or $\nabla = \lesssim$ and $p = \text{imprecise}$, for m being the approximation level of interest. Thus, a typical compatibility lemma for emulate looks like Lemma 12.

LEMMA 12 (COMPATIBILITY FOR λ EMULATION).

$$\begin{aligned} &\text{if } (m > n \text{ and } p = \text{precise}) \text{ or } (\nabla = \lesssim \text{ and } p = \text{imprecise}) \\ \text{then} \quad &\text{if } \text{toEmul}_{m;p}(\Gamma, x : \tau) \vdash t \nabla_n t : \text{EmulT}_{m;p;\tau}^{\text{fl}} \\ &\text{then } \text{toEmul}_{m;p}(\Gamma) \vdash \text{in-dn}_{m;\tau \rightarrow \tau'}^{\text{fl}}(\lambda x : \text{BtT}_{m;\tau}^{\text{fl}}. t) \nabla_n \lambda x : \tau. t : \text{EmulT}_{m;p;\tau \rightarrow \tau'}^{\text{fl}} \end{aligned}$$

The compatibility lemma for terms typed using type equality (Lemma 13) is the most interesting of these. The proof of this lemma is surprisingly simple because most of the heavy lifting is done by a corollary of Theorem 11, which proves that equivalent types have not only the same backtranslation type but also the same term relation.

LEMMA 13 (COMPATIBILITY LEMMA FOR EMULATION OF TYPE EQUALITY).

$$\begin{aligned} &\text{if } (m > n \text{ and } p = \text{precise}) \text{ or } (\nabla = \lesssim \text{ and } p = \text{imprecise}) \\ \text{then if } &\text{toEmul}_{m;p}^{\text{fe}}(\Gamma) \vdash t \nabla_n t : \text{EmulT}_{m;p;\tau}^{\text{fe}} \text{ and } \tau \stackrel{\text{def}}{=} \sigma \text{ then } \text{toEmul}_{m;p}^{\text{fe}}(\Gamma) \vdash t \nabla_n t : \text{EmulT}_{m;p;\sigma}^{\text{fe}} \end{aligned}$$

Corollary 1 (Equivalent types have the same term relation).

$$\text{if } \tau \stackrel{\text{def}}{=} \sigma \text{ then } \forall n. \mathcal{E} \llbracket \text{EmulT}_{n;p;\tau}^{\text{fe}} \rrbracket_{\nabla} = \mathcal{E} \llbracket \text{EmulT}_{n;p;\sigma}^{\text{fe}} \rrbracket_{\nabla}$$

Given a series of these kinds of compatibility lemmas, we can state that emulate is correct.

LEMMA 14 (EMULATE IS SEMANTICS-PRESERVING).

if $(m > n \text{ and } p = \text{precise}) \text{ or } (\nabla = \lesssim \text{ and } p = \text{imprecise}) \text{ and } \Gamma \vdash t : \tau$
 then $\text{toEmul}_{m;p}(\Gamma) \vdash \text{emulate}_m^{\text{fl}}(\Gamma \vdash t : \tau) \nabla_n t : \text{EmulT}_{m;p;\tau}^{\text{fl}}$

The key property we rely on for fully-abstract compilation though, is that emulation of contexts is correct (this relies on correctness of emulation for terms though).

LEMMA 15 (EMULATE IS SEMANTICS PRESERVING FOR CONTEXTS).

if $(m > n \text{ and } p = \text{precise}) \text{ or } (\nabla = \lesssim \text{ and } p = \text{imprecise}) \text{ and } \vdash \mathbb{C} : \Gamma', \tau' \rightarrow \Gamma, \tau$
 then $\vdash \text{emulate}_m^{\text{fl}}(\vdash \mathbb{C} : \Gamma', \tau' \rightarrow \Gamma, \tau) \nabla_n \mathbb{C} : \text{toEmul}_{m;p}(\Gamma'), \text{EmulT}_{m;p;\tau'}^{\text{fl}} \rightarrow \text{toEmul}_{m;p}(\Gamma), \text{EmulT}_{m;p;\tau}^{\text{fl}}$

4.2.2 Inject and Extract. As mentioned, the backtranslated target context must be a valid source context in order to be linked with a source term. Specifically, it must have a hole whose type is the compilation of some source type τ . Backtranslated terms, however, have backtranslation type $\text{BtT}_{n;\tau}^{\text{fe}}$, so we need to convert values of source type into values of backtranslation type (and back). To do this conversion we rely on functions $\text{inject}^{\text{fl}}$ and $\text{extract}^{\text{fl}}$ whose types and definitions are in Figure 12. Function $\text{inject}^{\text{fl}}$ takes a source value of type τ and converts it into “the same” value at the backtranslation type so that backtranslated terms can use that value. Since the backtranslation type is indexed by target types, we use function $\text{fxToIs}(\cdot)$ to generate the target type related to τ . Function $\text{extract}^{\text{fl}}$ does the dual and takes a value of backtranslation type and converts it into a type of some source type. These functions are defined mutually inductively in order to contravariantly convert function arguments to the appropriate type.

For values of the base type, these functions use the already introduced constructors and destructors for backtranslation type to perform their conversion. For pair and sum types, these functions operate recursively on the structure of the values they take in input. For arrow type, these functions convert the argument contravariantly before converting the result after the application of the function. When the size of the type is insufficient for these functions to behave as expected (i.e., when n is 0) it is sufficient for $\text{inject}^{\text{fl}}$ to return unit and for $\text{extract}^{\text{fl}}$ to just diverge.

Note that these functions are indexed by *source* types since they convert between them and the backtranslation type. Thus, while two of our compilers have the same source language (and therefore the same $\text{inject}/\text{extract}$), the third compiler has a different source language, with more types: $\mu\alpha.\tau$. Thus, for the third backtranslation, we have a different, extended version of $\text{inject}^{\text{IE}}/\text{extract}^{\text{IE}}$ that converts values of recursive types into values of backtranslation type and back. Additionally, the hole of the first two backtranslations cannot have a recursive type, since the source type for those backtranslations is λ^{fx} .

As for the emulation of terms, we prove that these functions are correct according to the logical relations. Terms that are related at a source type are related at backtranslation type after an $\text{inject}^{\text{fl}}$ while terms that are related at backtranslation type are related at source type after an $\text{extract}^{\text{fl}}$.

LEMMA 16 (INJECT AND EXTRACT ARE SEMANTICS PRESERVING).

If $(m \geq n \text{ and } p = \text{precise}) \text{ or } (\nabla = \lesssim \text{ and } p = \text{imprecise})$
 then if $\Gamma \vdash t \nabla_n t : \tau$ then $\Gamma \vdash \text{inject}_{m;\tau}^{\text{fl}} t \nabla_n t : \text{EmulT}_{m;p;\text{fxToIs}(\tau)}^{\text{fl}}$
 if $\Gamma \vdash t \nabla_n t : \text{EmulT}_{m;p;\text{fxToIs}(\tau)}^{\text{fl}}$ then $\Gamma \vdash \text{extract}_{m;\tau}^{\text{fl}} t \nabla_n t : \tau$

$\text{inject}_{n;\tau}^{\text{fl}} : \tau \rightarrow \text{BtT}_{n;\text{fxToIs}(\tau)}^{\text{fl}} \quad \text{and} \quad \text{extract}_{n;\tau}^{\text{fl}} : \text{BtT}_{n;\text{fxToIs}(\tau)}^{\text{fl}} \rightarrow \tau$

$$\begin{aligned}
\text{inject}_{0;\tau}^{\text{fl}} &= \lambda x : \tau. \text{unit} & \text{inject}_{n+1;\text{Unit}}^{\text{fl}} &= \lambda x : \text{Unit}. \text{inl } x & \text{inject}_{n+1;\text{Bool}}^{\text{fl}} &= \lambda x : \text{Bool}. \text{inl } x \\
\text{inject}_{n+1;\tau \rightarrow \tau'}^{\text{fl}} &= \lambda x : \tau \rightarrow \tau'. \text{inl } \lambda y : \text{BtT}_{n;\text{fxToIs}(\tau)}^{\text{fl}}. \text{inject}_{n;\tau'}^{\text{fl}} (x (\text{extract}_{n;\tau}^{\text{fl}} y)) \\
\text{inject}_{n+1;\tau \times \tau'}^{\text{fl}} &= \lambda x : \tau \times \tau'. \text{inl } \langle \text{inject}_{n;\tau}^{\text{fl}} (x.1), \text{inject}_{n;\tau'}^{\text{fl}} (x.2) \rangle \\
\text{inject}_{n+1;\tau \uplus \tau'}^{\text{fl}} &= \lambda x : \tau \uplus \tau'. \text{inl } \text{case } x \text{ of inl } x_1 \mapsto \text{inl } (\text{inject}_{n;\tau}^{\text{fl}} x_1) \mid \text{inr } x_2 \mapsto \text{inr } (\text{inject}_{n;\tau'}^{\text{fl}} x_2) \\
\text{extract}_{0;\tau}^{\text{fl}} &= \lambda x : \text{BtT}_{n;\text{fxToIs}(\tau)}^{\text{fl}}. \text{omega}_{\tau} \\
\text{extract}_{n+1;\text{Unit}}^{\text{fl}} &= \lambda x : \text{BtT}_{n+1;\text{Unit}}^{\text{fl}}. \text{Unit} \cdot \text{case}_{n+1;\text{Unit}}^{\text{fl}} \cdot \text{Unit } x \\
\text{extract}_{n+1;\text{Bool}}^{\text{fl}} &= \lambda x : \text{BtT}_{n+1;\text{Bool}}^{\text{fl}}. \text{Bool} \cdot \text{case}_{n+1;\text{Bool}}^{\text{fl}} \cdot \text{Bool } x \\
\text{extract}_{n+1;\tau \rightarrow \tau'}^{\text{fl}} &= \lambda x : \text{BtT}_{n+1;\text{fxToIs}(\tau \rightarrow \tau')}^{\text{fl}}. \lambda y : \tau. \text{extract}_{n;\tau'}^{\text{fl}} \left(\text{case}_{n+1;\text{fxToIs}(\tau \rightarrow \tau')}^{\text{fl}} \times (\text{inject}_{n;\tau}^{\text{fl}} y) \right) \\
\text{extract}_{n+1;\tau \times \tau'}^{\text{fl}} &= \lambda x : \text{BtT}_{n+1;\text{fxToIs}(\tau \times \tau')}^{\text{fl}}. \left\langle \text{extract}_{n;\tau}^{\text{fl}} \left(\text{case}_{n+1;\text{fxToIs}(\tau \times \tau')}^{\text{fl}} x.1 \right), \right. \\
&\quad \left. \text{extract}_{n;\tau'}^{\text{fl}} \left(\text{case}_{n+1;\text{fxToIs}(\tau \times \tau')}^{\text{fl}} x.2 \right) \right\rangle \\
&\quad \text{case } \left(\text{case}_{n+1;\text{fxToIs}(\tau \uplus \tau')}^{\text{fl}} x \right) \text{ of} \\
\text{extract}_{n+1;\tau \uplus \tau'}^{\text{fl}} &= \lambda x : \text{BtT}_{n+1;\text{fxToIs}(\tau \uplus \tau')}^{\text{fl}}. \left| \begin{array}{l} \text{inl } x_1 \mapsto \text{inl } \text{extract}_{n;\text{fxToIs}(\tau)}^{\text{fl}} x_1 \\ \text{inr } x_2 \mapsto \text{inr } \text{extract}_{n;\text{fxToIs}(\tau')}^{\text{fl}} x_2 \end{array} \right.
\end{aligned}$$

$\text{inject}_{n;\tau}^{\text{fe}} \stackrel{\text{def}}{=} \text{as above}$

$\text{extract}_{n;\tau}^{\text{fe}} \stackrel{\text{def}}{=} \text{as above}$

$$\begin{aligned}
\text{inject}_{n+1;\mu\alpha.\tau}^{\text{IE}} &= \lambda x : \mu\alpha.\tau. \text{inject}_{n+1;\tau[\mu\alpha.\tau/\alpha]}^{\text{IE}} \quad (\text{unfold } \mu\alpha.\tau \ x) \\
\text{extract}_{n+1;\mu\alpha.\tau}^{\text{IE}} &= \lambda x : \text{BtT}_{n+1;\text{isToEq}(\mu\alpha.\tau)}^{\text{IE}}. \text{extract}_{n+1;\mu\alpha.\tau}^{\text{IE}} \text{ fold } \mu\alpha.\tau \ (\text{case}_{n+1;\text{isToEq}(\mu\alpha.\tau)}^{\text{IE}} x) \\
&\quad \text{omitted cases are as above}
\end{aligned}$$

Fig. 12. Definition of the `inject` and `extract` functions.

4.2.3 The Backtranslations. The backtranslation of a target context based on its type derivation is defined as follows by relying on both `emulatefl` and `injectfl`. All three backtranslations follow exactly the same pattern and enjoy the same properties. As already shown, the only interesting changes are in the sub-parts of the backtranslation (e.g., in the different definitions of `inject/extract`). Thus, we only show the backtranslation from λ_I^μ to λ^{fx} and we state properties only for this one.

Definition 3 (Approximate backtranslation for λ_I^μ contexts into λ^{fx}).

$$\ll \mathbb{C}, n \gg_{\lambda^{\text{fx}}}^{\lambda_I^\mu} \stackrel{\text{def}}{=} \text{emulate}_n^{\text{fl}} \left(\vdash \mathbb{C} : \Gamma, \ll \tau \gg_{\lambda_I^\mu}^{\lambda^{\text{fx}}} \rightarrow \Gamma', \tau' \right) \left[\text{inject}_{n;\tau}^{\text{fl}} \cdot \right] \quad (\text{provided } \vdash \mathbb{C} : \Gamma, \ll \tau \gg_{\lambda_I^\mu}^{\lambda^{\text{fx}}} \rightarrow \Gamma', \tau')$$

As for the compiler from λ^{fx} to λ_E^μ , we can derive the backtranslation from λ_E^μ to λ^{fx} by composing the backtranslations through λ_I^μ . Thus, $\ll t \gg_{\lambda^{\text{fx}}}^{\lambda_E^\mu} = \ll \ll t \gg_{\lambda_I^\mu}^{\lambda_E^\mu} \gg_{\lambda^{\text{fx}}}^{\lambda_I^\mu}$. Interestingly, this means that the type of λ_E^μ terms backtranslated into λ^{fx} is the same as the one for λ_E^μ terms backtranslated into λ_I^μ , i.e., the case for `BtTfe` for $\mu\alpha.\tau$ should not lose precision (as shown in Figure 5). Notice that the

first backtranslation ($\langle\langle \cdot \rangle\rangle_{\lambda_1^\mu}^{\lambda_E^\mu}$) directs this, since $\mathbf{BtT}^{\text{IE}}$ is simply a collection of $\hat{\tau} \cup \hat{\tau}'$ pseudotypes, the second backtranslation ($\langle\langle \cdot \rangle\rangle_{\lambda_{fx}^\mu}^{\lambda_1^\mu}$) simply relies on the case for $\mathbf{BtT}_{n;\tau \cup \tau'}^{\text{fl}}$.

Using the same approach for the correctness of emulate, we can state that the backtranslations are correct. For simplicity, we provide a visual representation of this proof in Figure 13 (adapted from the work of Devriese et al. [2016] to our setting). All of the infrastructure used by the backtranslation (i.e., $\text{inject}^{\text{fl}}$ / $\text{extract}^{\text{fl}}$ and the $\mathbf{BtT}^{\text{fl}}$ helpers) have correctness lemmas that follow the same structure of the one for $\text{emulate}^{\text{fl}}(\cdot)$. Specifically, they relate terms at $\mathbf{EmulT}^{\text{fl}}$, they transform target environments into source ones via function $\text{toEmul}(\cdot)$ and they have a condition on the different directions of the approximation (the first line in Lemmas 12 to 15).

Fig. 13. Diagram representing the relatedness between different bits of the backtranslation and of the compiler.

LEMMA 17 (CORRECTNESS OF $\langle\langle \cdot \rangle\rangle_{\lambda_{fx}^\mu}^{\lambda_1^\mu}$).

If $(m \geq n \text{ and } p = \text{precise}) \text{ or } (\nabla = \lesssim \text{ and } p = \text{imprecise})$

then if $\vdash \mathbf{C} : \emptyset, \llbracket \tau \rrbracket_{\lambda_1^\mu}^{\lambda_{fx}^\mu} \rightarrow \emptyset, \tau$ and $\emptyset \vdash t \nabla_n t : \tau$ then $\emptyset \vdash \langle\langle \mathbf{C}, m \rangle\rangle_{\lambda_{fx}^\mu}^{\lambda_1^\mu} [t] \nabla_n \mathbf{C}[t] : \mathbf{EmulT}_{m;p;\tau}^{\text{fl}}$

With correctness of the backtranslation we can prove the preservation direction of fully-abstract compilation for all compilers, following the proof structure of Figure 4.

Theorem 18 ($\llbracket \cdot \rrbracket_{\lambda_1^\mu}^{\lambda_{fx}^\mu}$ preserves equivalence).

If $\emptyset \vdash t_1 \simeq_{\text{ctx}} t_2 : \tau$ then $\emptyset \vdash \llbracket t_1 \rrbracket_{\lambda_1^\mu}^{\lambda_{fx}^\mu} \simeq_{\text{ctx}} \llbracket t_2 \rrbracket_{\lambda_1^\mu}^{\lambda_{fx}^\mu} : \llbracket \tau \rrbracket_{\lambda_1^\mu}^{\lambda_{fx}^\mu}$

PROOF. Take $\mathbf{C}$ such that $\vdash \mathbf{C} : \emptyset, \llbracket \tau \rrbracket_{\lambda_1^\mu}^{\lambda_{fx}^\mu} \rightarrow \emptyset, \tau$. We need to prove that $\mathbf{C} [\llbracket t_1 \rrbracket_{\lambda_1^\mu}^{\lambda_{fx}^\mu}] \Downarrow \iff \mathbf{C} [\llbracket t_2 \rrbracket_{\lambda_1^\mu}^{\lambda_{fx}^\mu}] \Downarrow$. By symmetry, we prove only that if $\mathbf{C} [\llbracket t_1 \rrbracket_{\lambda_1^\mu}^{\lambda_{fx}^\mu}] \Downarrow$ then $\mathbf{C} [\llbracket t_2 \rrbracket_{\lambda_1^\mu}^{\lambda_{fx}^\mu}] \Downarrow$ (HPTT). Take n strictly larger than the steps needed for $\mathbf{C} [\llbracket t_1 \rrbracket_{\lambda_1^\mu}^{\lambda_{fx}^\mu}] \Downarrow$. By Lemma 4 ($\llbracket \cdot \rrbracket_{\lambda_1^\mu}^{\lambda_{fx}^\mu}$ is semantics preserving) we have $\emptyset \vdash t_1 \nabla_n \llbracket t_1 \rrbracket_{\lambda_1^\mu}^{\lambda_{fx}^\mu} : \tau$. Take $m = n$, so we have $(m \geq n \text{ and } p = \text{precise})$ and therefore $(\nabla = \gtrsim)$. By Lemma 17 (Correctness of $\langle\langle \cdot \rangle\rangle_{\lambda_{fx}^\mu}^{\lambda_1^\mu}$) we have $\emptyset \vdash \langle\langle \mathbf{C}, m \rangle\rangle_{\lambda_{fx}^\mu}^{\lambda_1^\mu} [t_1] \gtrsim_n \mathbf{C} [\llbracket t_1 \rrbracket_{\lambda_1^\mu}^{\lambda_{fx}^\mu}] : \mathbf{EmulT}_{m;p;\tau}^{\text{fl}}$. By Lemma 2 (Adequacy for $\gtrsim$) and HPTT we have: $\langle\langle \mathbf{C}, m \rangle\rangle_{\lambda_{fx}^\mu}^{\lambda_1^\mu} [t_1] \Downarrow$, which by source contextual equivalence gives us $\langle\langle \mathbf{C}, m \rangle\rangle_{\lambda_{fx}^\mu}^{\lambda_1^\mu} [t_2] \Downarrow$ (HPTS2). Given n' the number of steps for HPTS2, by Lemma 4 ($\llbracket \cdot \rrbracket_{\lambda_1^\mu}^{\lambda_{fx}^\mu}$ is semantics preserving) we have: $\emptyset \vdash t_2 \nabla_{n'} \llbracket t_2 \rrbracket_{\lambda_1^\mu}^{\lambda_{fx}^\mu} : \tau$. So by definition: $\emptyset \vdash t_2 \lesssim_{n'} \llbracket t_2 \rrbracket_{\lambda_1^\mu}^{\lambda_{fx}^\mu} : \tau$. By Lemma 17 (Correctness of $\langle\langle \cdot \rangle\rangle_{\lambda_{fx}^\mu}^{\lambda_1^\mu}$) (with $n = n'$, $p = \text{imprecise}$ and

$\nabla = \lesssim$) we can conclude $\emptyset \vdash \llbracket \mathbf{C}, \mathbf{m} \rrbracket_{\lambda_1^\mu}^{t_2} \lesssim_n \mathbf{C} \left[\llbracket t_2 \rrbracket_{\lambda_1^\mu}^{\lambda^{\text{fx}}} \right] : \text{EmulT}_{m;p;\tau}^{\text{fl}}$. By Lemma 2 (Adequacy for $\approx$) for $\lesssim$ with HPTS2 we conclude the thesis. $\square$

Theorem 19 ($\llbracket \cdot \rrbracket_{\lambda_1^\mu}^{\lambda^{\text{fx}}}$ preserves equivalence).

If $\emptyset \vdash t_1 \approx_{\text{ctx}} t_2 : \tau$ then $\emptyset \vdash \llbracket t_1 \rrbracket_{\lambda_1^\mu}^{\lambda^{\text{fx}}} \approx_{\text{ctx}} \llbracket t_2 \rrbracket_{\lambda_1^\mu}^{\lambda^{\text{fx}}} : \llbracket \tau \rrbracket_{\lambda_1^\mu}^{\lambda^{\text{fx}}}$

Theorem 20 ($\llbracket \cdot \rrbracket_{\lambda_1^\mu}^{\lambda^{\text{fx}}}$ preserves equivalence).

If $\emptyset \vdash t_1 \approx_{\text{ctx}} t_2 : \tau$ then $\emptyset \vdash \llbracket t_1 \rrbracket_{\lambda_1^\mu}^{\lambda^{\text{fx}}} \approx_{\text{ctx}} \llbracket t_2 \rrbracket_{\lambda_1^\mu}^{\lambda^{\text{fx}}} : \llbracket \tau \rrbracket_{\lambda_1^\mu}^{\lambda^{\text{fx}}}$

4.3 Full Abstraction for the Three Compilers

With the two directions of fully-abstract compilation already proved, we can easily show that all three compilers are fully abstract. As before, full abstraction of $\llbracket \cdot \rrbracket_{\lambda_1^\mu}^{\lambda^{\text{fx}}}$ trivially follows from composing full abstraction for the other two compilers.

Theorem 21 ($\llbracket \cdot \rrbracket_{\lambda_1^\mu}^{\lambda^{\text{fx}}}$ is fully abstract). $\emptyset \vdash t_1 \approx_{\text{ctx}} t_2 : \tau \iff \emptyset \vdash \llbracket t_1 \rrbracket_{\lambda_1^\mu}^{\lambda^{\text{fx}}} \approx_{\text{ctx}} \llbracket t_2 \rrbracket_{\lambda_1^\mu}^{\lambda^{\text{fx}}} : \llbracket \tau \rrbracket_{\lambda_1^\mu}^{\lambda^{\text{fx}}}$

Theorem 22 ($\llbracket \cdot \rrbracket_{\lambda_1^\mu}^{\lambda^{\text{fx}}}$ is fully abstract). $\emptyset \vdash t_1 \approx_{\text{ctx}} t_2 : \tau \iff \emptyset \vdash \llbracket t_1 \rrbracket_{\lambda_1^\mu}^{\lambda^{\text{fx}}} \approx_{\text{ctx}} \llbracket t_2 \rrbracket_{\lambda_1^\mu}^{\lambda^{\text{fx}}} : \llbracket \tau \rrbracket_{\lambda_1^\mu}^{\lambda^{\text{fx}}}$

Theorem 23 ($\llbracket \cdot \rrbracket_{\lambda_1^\mu}^{\lambda^{\text{fx}}}$ is fully abstract). $\emptyset \vdash t_1 \approx_{\text{ctx}} t_2 : \tau \iff \emptyset \vdash \llbracket t_1 \rrbracket_{\lambda_1^\mu}^{\lambda^{\text{fx}}} \approx_{\text{ctx}} \llbracket t_2 \rrbracket_{\lambda_1^\mu}^{\lambda^{\text{fx}}} : \llbracket \tau \rrbracket_{\lambda_1^\mu}^{\lambda^{\text{fx}}}$

5 RELATED WORK

Two alternative formulations of equi-recursive types exist: one based on an inductive type equality (which we dub λ_{Ei}^μ in this section) and one based on a weak type equality (which we dub λ_{Es}^μ).⁶ λ_{Ei}^μ defines an equality relation on types ($\approx$) that, unlike ours, is inductively defined [Abadi and Fiore 1996]. Types are equal if they are the same (Rules Eq-type-Base and Eq-type-Var), when their subparts are equal (Rules Eq-type-Bi and Eq-type-Mu) or when one is the unfolding of the other (Rule Eq-type-Unfold). To keep track of type variables, typing equality is defined with respect to an environment $\Delta ::= \emptyset \mid \Delta; \alpha$.

$\tau \approx \sigma$			
(Eq-type-Symmetric)	(Eq-type-Transitive)	(Eq-type-Bi)	(Eq-type-Base)
$\frac{\Delta \vdash \tau' \approx \tau}{\Delta \vdash \tau \approx \tau'}$	$\frac{\Delta \vdash \tau \approx \tau'' \quad \Delta \vdash \tau'' \approx \tau'}{\Delta \vdash \tau \approx \tau'}$	$\frac{\star \in \{\rightarrow, \times, \cup\} \quad \Delta \vdash \tau_1 \approx \tau'_1 \quad \Delta \vdash \tau_2 \approx \tau'_2}{\Delta \vdash \tau_1 \star \tau_2 \approx \tau'_1 \star \tau'_2}$	$\frac{i = \text{Unit} \vee i = \text{Bool}}{\Delta \vdash i \approx i}$
(Eq-type-Var)	(Eq-type-Mu)	(Eq-type-Unfold)	
$\frac{\alpha \in \Delta}{\Delta \vdash \alpha \approx \alpha}$	$\frac{\Delta, \alpha \vdash \tau \approx \tau'}{\Delta \vdash \mu \alpha. \tau \approx \mu \alpha. \tau'}$	$\frac{\Delta \vdash \tau[\mu \alpha. \tau / \alpha] \approx \tau'}{\Delta \vdash \mu \alpha. \tau \approx \tau'}$	

Cai et al. [2016] explain that this notion of type equality is strictly weaker than the coinductive one we have used. For example, they mention two type equalities that do not hold in λ_{Ei}^μ :

$$\emptyset \vdash \mu \alpha. \alpha \rightarrow \text{Unit} \neq \mu \alpha. (\alpha \rightarrow \text{Unit}) \rightarrow \text{Unit} \quad \emptyset \vdash \mu \alpha. \mu \beta. \alpha \rightarrow \beta \neq \mu \alpha. \alpha \rightarrow \alpha$$

To understand why these equalities do not hold in the inductive formulation, consider that no amount of unfolding of a recursive type μs will ever produce recursive types with a different body.

⁶ We typeset these languages in a **green**, **verbatim** font, though they appear in this section only.

λ_{Es}^μ instead enforces that just a recursive type and its unfolding are equivalent [Ahmed 2004; Appel and McAllester 2001; MacQueen et al. 1986; Urzyczyn 1995]. This leads to more compact typing rules and it does not require a type equivalence relation, effectively this is like λ_{I}^μ but without **fold/unfold** annotations.

The main difference is that in this last variant, unfoldings can only happen at the top-level of a type of a term (i.e., when terms are of a recursive type themselves). In both λ_{Ei}^μ and in our coinductive variant λ_E^μ , unfoldings can also happen inside the types. For example, types such as $(\mu\alpha. B \uplus \alpha) \rightarrow B$ and $(B \uplus (\mu\alpha. B \uplus \alpha)) \rightarrow B$ are not equivalent in this last variant, because we can unfold $\mu\alpha. B \uplus \alpha$ to $(B \uplus (\mu\alpha. B \uplus \alpha))$ inside the domain of the function type. These types are however equivalent in λ_{Ei}^μ and in λ_E^μ .

Since terms of λ_{Ei}^μ (or λ_{Es}^μ) can be typed in λ_E^μ and their semantics do not vary, our results show that all these different formulations of equi-recursive types are equally expressive. Since the approximate backtranslation is needed to deal with the coinductive derivations of λ_E^μ , we believe that a precise backtranslation akin to that of New et al. [2016] can be used to prove full abstraction for the compiler from λ_{I}^μ to λ_{Ei}^μ . We leave investigating this for future work.

As mentioned in Section 1, the closest work to ours is that of Abadi and Fiore [1996]. Like us, they study the relation between iso- and equi-recursive types and prove that any term typed λ_{I}^μ can be typed in λ_{Ei}^μ and vice versa. For the backward direction, they insert cast functions which appropriately insert **fold** and **unfold** annotations to make terms typecheck. Additionally, they use a logic to prove that the terms with the casts are equivalent to the original, but the logic does not come with a soundness proof. Abadi and Fiore do not connect their results to the operational semantics in any way, unlike ours, and their results cannot be used to derive fully-abstract compilation, as they relate one term and its compilation, not two terms and their compilation. Finally, it is not clear if Abadi and Fiore’s Theorem 6.8 can be interpreted to imply any form of equi-expressiveness of the two languages. In fact, what Abadi and Fiore prove is that an equi-recursive term is equal to a back-translated term under a certain equality that is (conjectured to be) almost (but not entirely) sound for observational equivalence in equi-recursive contexts. On the other hand, in our setting, the interaction of the same programs with arbitrary contexts provides a measure on the relative expressiveness of those contexts when interacting with the given programs. This difference is key to make claims about the relative expressive power of languages, as we make.

Fully-abstract compilation derived from fully-abstract semantics models [Milner 1977], and it has been initially devised to study the relative expressive power of programming languages [Felleisen 1991; Gorla and Nestmann 2016; Mitchell 1993].⁷ Fully-abstract compilation has been widely used to compare process algebras and their relative expressiveness, as surveyed by Parrow [2008]. Additionally, researchers have argued that fully-abstract compilation is a feasible criterion for secure compilation [Abadi 1998; Kennedy 2006], as surveyed by Patrignani et al. [2019].

Proofs of fully-abstract compilation are notoriously complex and thus a large amount of work exists in devising proof techniques for it. Most of these proof techniques require a form of back-translation [Ahmed and Blume 2008, 2011; Bowman and Ahmed 2015]. Precise backtranslations generate source contexts that reproduce the behaviour of the target context faithfully, without any approximation [New et al. 2016; Van Strydonck et al. 2019]. Approximate backtranslations, instead, generate source contexts that reproduce that behaviour up to a certain number of steps. The approximate backtranslation proof technique we use was conjectured by Schmidt-Schauß et al. [2015] and was used by Devriese et al. [2017] to prove full abstraction for a compiler from λ^{fx} to the

⁷ Not all these works use the term “fully-abstract compilation” but their intuition is the same.

untyped lambda calculus (λ^u). Unlike these works, we deal with a family of backtranslation types that is indexed by target types. Additionally, our compilers do not perform dynamic typechecks; they are simply the canonical translation of a term in the source language into the target. Finally, we remark that our results cannot be derived from Devriese et al. [2016] since the languages in that paper have no recursive types.

Interestingly, our current result can be seen as factoring out the first phase of Devriese et al. [2016]’s compiler; their result could be seen as composing one of our current results with a second fully abstract compiler from λ_1^μ to λ^u , which takes care of dynamic type enforcement. The full abstraction proof for this second compiler could be a lot simpler with recursive types in the source language, as it would no longer require an approximate backtranslation. In fact, we believe that reusable sub-results could be factored out from other full abstraction results in the literature too. For example, we conjecture that one could separate closure conversion from purity enforcement in New et al. [2016]’s compiler, or separate contract enforcement from universal contract erasure in Van Strydonck et al. [2019]’s compiler. We hope our experience can inspire other researchers to pay more attention to such factoring opportunities and strive to minimize compiler phases. In other words, we believe the community could benefit from using a nanopass secure compilation mindset, in the spirit of nanopass compilation [Sarkar et al. 2004]. Even computationally-trivial nanopasses like ours can be useful as they enrich the power of contexts and simplify secure compilation proofs further downstream.

6 CONCLUSION

This paper demonstrates that the simply typed lambda calculus with iso- and equi-recursive types has the same expressive power. To do so, it presented three fully-abstract compilers in order to reason about iso- and equi-recursively typed terms interacting over a simply-typed interface and a recursively-typed one. The first compiler translates from a simply-typed lambda calculus with a fixpoint operator (λ^{fx}) to a simply-typed lambda calculus with iso-recursive types (λ_1^μ). The second compiler translates from λ^{fx} to a simply-typed lambda calculus with coinductive equi-recursive types (λ_E^μ). These two compilers demonstrate the same expressive power of iso- and equi-recursive types on a simply-typed interface. The third compiler translates from λ_1^μ to λ_E^μ , demonstrating equal expressiveness of iso- and equi-recursive types on a recursively-typed interface. All fully-abstract compilation proofs rely on a novel adaptation of the approximate backtranslation proof technique that works with families of target types-indexed backtranslation type.

ACKNOWLEDGMENTS

The authors thank the anonymous reviewers for detailed feedback on an earlier draft as well as Phil Wadler for interesting comments and suggestions. This work was partially supported: by the Office of Naval Research for support through grant N00014-18-1-2620, Accountable Protocol Customization, by the German Federal Ministry of Education and Research (BMBF) through funding for the CISPA-Stanford Center for Cybersecurity (FKZ: 13N1S0762), by the Air Force Office of Scientific Research under award number FA9550-21-1-0054, and by the Fund for Scientific Research - Flanders (FWO).

REFERENCES

- Martín Abadi. 1998. Protection in Programming-Language Translations. In *ICALP’98*. 868–883.
- Martin Abadi and Marcelo P. Fiore. 1996. Syntactic Considerations on Recursive Types. In *Proceedings of the 11th Annual IEEE Symposium on Logic in Computer Science (LICS ’96)*. IEEE Computer Society, Washington, DC, USA, 242–.
- Amal Ahmed. 2004. *Semantics of Types for Mutable State*. Ph.D. Dissertation. Princeton University.

- Amal Ahmed and Matthias Blume. 2008. Typed Closure Conversion Preserves Observational Equivalence. In *International Conference on Functional Programming*. ACM, 157–168.
- Amal Ahmed and Matthias Blume. 2011. An Equivalence-Preserving CPS Translation via Multi-Language Semantics. In *Proceedings of the 16th ACM SIGPLAN International Conference on Functional Programming (Tokyo, Japan) (ICFP '11)*. ACM, 431–444.
- Andrew W. Appel and David McAllester. 2001. An Indexed Model of Recursive Types for Foundational Proof-carrying Code. *ACM Trans. Program. Lang. Syst.* 23, 5 (Sept. 2001), 657–683.
- Nick Benton and Chung-Kil Hur. 2009. Biorthogonality, step-indexing and compiler correctness. *SIGPLAN Not.* 44, 97–108.
- William J. Bowman and Amal Ahmed. 2015. Noninterference for free. In *ICFP*. ACM.
- Yufei Cai, Paolo G. Giarrusso, and Klaus Ostermann. 2016. System F-omega with Equirecursive Types for Datatype-generic Programming. *SIGPLAN Not.* 51, 1 (Jan. 2016), 30–43.
- Dominique Devriese, Marco Patrignani, and Frank Piessens. 2016. Fully-abstract Compilation by Approximate Back-translation. In *Principles of Programming Languages*. 164–177.
- Dominique Devriese, Marco Patrignani, Frank Piessens, and Steven Keuchel. 2017. Modular, Fully-abstract Compilation by Approximate Back-translation. *Logical Methods in Computer Science* Volume 13, Issue 4 (Oct. 2017).
- Matthias Felleisen. 1991. On the Expressive Power of Programming Languages. In *Selected Papers from the Symposium on 3rd European Symposium on Programming (ESOP '90)*. Elsevier North-Holland, Inc., New York, NY, USA, 35–75.
- Cedric Fournet, Nikhil Swamy, Juan Chen, Pierre-Evariste Dagand, Pierre-Yves Strub, and Benjamin Livshits. 2013. Fully Abstract Compilation to JavaScript. In *Principles of Programming Languages*. ACM, 371–384.
- M. Gordon, R. Milner, and C. P. Wadsworth. 1979. *Edinburgh LCF: A Mechanized Logic of Computation*. Springer-Verlag, Berlin Heidelberg. <https://doi.org/10.1007/3-540-09724-4>
- Daniele Gorla and Uwe Nestmann. 2016. Full abstraction for expressiveness: history, myths and facts. *Mathematical Structures in Computer Science* 26, 4 (2016), 639–654.
- Robert Harper and John C. Mitchell. 1993. On the Type Structure of Standard ML. *ACM Transactions on Programming Languages and Systems* 15, 2 (April 1993), 211–252. <https://doi.org/10.1145/169701.169696>
- Chung-Kil Hur and Derek Dreyer. 2011. A Kripke Logical Relation Between ML and Assembly. In *Principles of Programming Languages*. 133–146.
- Hyeonseung Im, Keiko Nakata, and Sungwoo Park. 2013. Contractive Signatures with Recursive Types, Type Parameters, and Abstract Types. In *Automata, Languages, and Programming*, Fedor V. Fomin, Rūsiņš Freivalds, Marta Kwiatkowska, and David Peleg (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 299–311.
- Andrew Kennedy. 2006. Securing the .NET Programming Model. *Theoretical Computer Science* 364 (2006), 311–317.
- David MacQueen, Gordon Plotkin, and Ravi Sethi. 1984. An Ideal Model for Recursive Polymorphic Types. In *Proceedings of the 11th ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages (Salt Lake City, Utah, USA) (POPL '84)*. Association for Computing Machinery, New York, NY, USA, 165–174. <https://doi.org/10.1145/800017.800528>
- David MacQueen, Gordon Plotkin, and Ravi Sethi. 1986. An ideal model for recursive polymorphic types. *Information and Control* 71, 1 (1986), 95 – 130.
- Robin Milner. 1977. Fully abstract models of typed λ -calculi. *Theoretical Computer Science* 4, 1 (1977), 1 – 22.
- John C. Mitchell. 1993. On abstraction and the expressive power of programming languages. *Science of Computer Programming* 21, 2 (1993), 141 – 163.
- James H. Morris. 1968. *Lambda-Calculus Models of Programming Languages*. Ph.D. Dissertation. Massachusetts Institute of Technology.
- Max S. New, William J. Bowman, and Amal Ahmed. 2016. Fully Abstract Compilation via Universal Embedding. In *International Conference on Functional Programming*. ACM, 103–116.
- Joachim Parrow. 2008. Expressiveness of Process Algebras. *Elec. Not. Theo. Comp. Sci.* 209, 0 (2008), 173 – 186.
- Marco Patrignani. 2020. Why Should Anyone use Colours? or, Syntax Highlighting Beyond Code Snippets. CoRR abs/2001.11334.
- Marco Patrignani, Pieter Agten, Raoul Strackx, Bart Jacobs, Dave Clarke, and Frank Piessens. 2015. Secure Compilation to Protected Module Architectures. *ACM Trans. Program. Lang. Syst.* 37, Article 6 (April 2015), 6:1–6:50 pages.
- Marco Patrignani, Amal Ahmed, and Dave Clarke. 2019. Formal Approaches to Secure Compilation A Survey of Fully Abstract Compilation and Related Work. *ACM Comput. Surv.* 51, 6, Article 125 (Jan. 2019), 36 pages.
- Marco Patrignani, Eric Mark Martin, and Dominique Devriese. 2020. On the Semantic Expressiveness of Recursive Types. arXiv:2010.10859 [cs.PL]
- Benjamin Pierce. 2002. *Types and Programming Languages*. MIT Press.
- Gordon D. Plotkin. 1977. LCF Considered as a Programming Language. *Theoretical Computer Science* 5 (1977), 223–255.
- Dipanwita Sarkar, Oscar Waddell, and R. Kent Dybvig. 2004. A Nanopass Infrastructure for Compiler Education. *ACM SIGPLAN Notices* 39, 9 (Sept. 2004), 201–212. <https://doi.org/10.1145/1016848.1016878>

- Manfred Schmidt-Schauß, David Sabel, Joachim Niehren, and Jan Schwinghammer. 2015. Observational program calculi and the correctness of translations. *Theoretical Computer Science* 577 (2015), 98 – 124.
- Lau Skorstengaard, Dominique Devriese, and Lars Birkedal. 2019. StkTokens: Enforcing Well-Bracketed Control Flow and Stack Encapsulation Using Linear Capabilities. *Proc. ACM Program. Lang.* 3, POPL (Jan. 2019), 19:1–19:28.
- Pawel Urzyczyn. 1995. Positive Recursive Type Assignment. In *Proceedings of the 20th International Symposium on Mathematical Foundations of Computer Science (MFCS '95)*. Springer-Verlag, Berlin, Heidelberg, 382–391.
- Thomas Van Strydonck, Frank Piessens, and Dominique Devriese. 2019. Linear Capabilities for Fully Abstract Compilation of Separation-Logic-Verified Code. *Proc. ACM Program. Lang.* ICFP (2019).