

Fully-Abstract Compilation by Approximate Back-Translation

Dominique Devriese Marco Patrignani Frank Piessens

iMinds-Distrinet, Computer Science dept., KU Leuven, Belgium

first.last @ cs.kuleuven.be

Abstract

A compiler is *fully-abstract* if the compilation from source language programs to target language programs preserves and reflects behavioural equivalence. Such compilers have important security benefits, as they limit the power of an attacker interacting with the program in the target language to that of an attacker interacting with the program in the source language. Proving compiler full-abstractness is, however, rather complicated. A common proof technique is based on the *back-translation* of target-level program contexts to behaviourally-equivalent source-level contexts. However, constructing such a back-translation is problematic when the source language is not strong enough to embed an encoding of the target language. For instance, when compiling from the simply-typed λ -calculus (λ^τ) to the untyped λ -calculus (λ^u), the lack of recursive types in λ^τ prevents such a back-translation.

We propose a general and elegant solution for this problem. The key insight is that it suffices to construct an *approximate* back-translation. The approximation is only accurate up to a certain number of steps and conservative beyond that, in the sense that the context generated by the back-translation may diverge when the original would not, but not vice versa. Based on this insight, we describe a general proof technique for proving compiler full-abstractness and demonstrate it on a compiler from λ^τ to λ^u . The proof uses asymmetric cross-language logical relations and makes innovative use of step-indexing to express the relation between a context and its approximate back-translation. We believe this proof technique can scale to challenging settings and enable simpler, more scalable proofs of compiler full-abstractness.

Keywords Fully-abstract compilation, cross-language logical relations, step-indexing, compiler security, secure compilation

We use **green** and **pink** to typeset resp. source and target language terms; we recommend to print this paper in colour for maximum clarity.

1. Introduction

A compiler is *fully-abstract* if the compilation from source language programs to target language programs preserves and reflects behavioural equivalence [1, 12]. Such compilers have important security benefits. It is often realistic to assume that attackers can interact with a program in the target language, and depending on the target language this can enable attacks such as improper stack manipulation, breaking control flow guarantees, reading from or writ-




Figure 1. Proving one half of full-abstractness: compiler correctness. Only one direction of this half is presented (\Rightarrow), the other one follows by symmetry.

ing to private memory of other components, inspecting or modifying the implementation of a function etc [1, 2, 11, 16, 24]. A fully-abstract compiler is sufficiently defensive to rule out such attacks: the power of an attacker interacting with the program in the target language is limited to attacks that could also be performed by an attacker interacting with the program in the source language.

Formally, we model a compiler as a function $\llbracket \cdot \rrbracket$ that maps source language terms t to target language terms $\llbracket t \rrbracket$. Elements of the source language are typeset in a **green, bold** font, while elements of the target language are typeset in a **pink, sans-serif** font. Roughly, the compiler is fully-abstract, if for any two source language terms t_1 and t_2 , we have that they are behaviourally equivalent ($t_1 \approx_{ctx} t_2$) if and only if their compiled counterparts are behaviourally equivalent ($\llbracket t_1 \rrbracket \approx_{ctx} \llbracket t_2 \rrbracket$) [1]. The notion of behavioural equivalence used here is the canonical notion of contextual equivalence: two terms are equivalent if they behave the same when plugged into any valid context. Specifically, we take contextual equivalence to be equi-termination: $t \approx_{ctx} t' \stackrel{\text{def}}{=} \forall \mathcal{C}, \mathcal{C}[t] \Downarrow \iff \mathcal{C}[t'] \Downarrow$. The universal quantification over contexts \mathcal{C} ensures that the results produced by t and t' are the same [9, 27].

The full-abstractness property can be split into two parts: the right-to-left implication and the left-to-right implication, which we call compiler *correctness* and compiler *security* respectively.

Compiler correctness ($t_1 \approx_{ctx} t_2 \Leftarrow \llbracket t_1 \rrbracket \approx_{ctx} \llbracket t_2 \rrbracket$) requires that if the compiler produces equivalent target programs, then the source programs must have been equivalent. Intuitively, it is clear that if programs with different source language behaviour become equivalent after compilation, the compiler must have incorrectly compiled at least one of them.

We build on cross-language logical relations: a technique that has recently been proposed for proving compiler correctness [6, 7, 13]. The general idea of this approach is depicted in Fig. 1 (purposely ignoring language-specific things such as the types of the terms involved). The proof starts from the knowledge that $\llbracket t_1 \rrbracket \approx_{ctx} \llbracket t_2 \rrbracket$ and needs to prove that $t_1 \approx_{ctx} t_2$. That is, for an arbitrary valid context \mathcal{C} , it shows that $\mathcal{C}[t_1] \Downarrow$ if and only if $\mathcal{C}[t_2] \Downarrow$. By symmetry, it suffices to show that $\mathcal{C}[t_1] \Downarrow \Rightarrow \mathcal{C}[t_2] \Downarrow$.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

POPL '16 Month d-d, 2016, City, ST, Country
Copyright © 2016 ACM [to be supplied]...\$15.00




Figure 2. Proving the other half of full-abstraction: compiler security.

The idea of the approach is to define a cross-language logical relation $t \approx t$ that expresses when a compiled term t behaves as a target-level version of source-level term t . This logical relation is not compiler-specific: it should be understood as a specification of a target-level calling convention rather than a precise compiler result. If we can then prove that any term is logically related to its compilation ($t \approx \llbracket t \rrbracket$), and that the same result holds for contexts ($\mathcal{C} \approx \llbracket \mathcal{C} \rrbracket$), then compiler correctness follows. Starting from $t_1 \approx \llbracket t_1 \rrbracket$ and $t_2 \approx \llbracket t_2 \rrbracket$ and $\mathcal{C} \approx \llbracket \mathcal{C} \rrbracket$, the proof uses the inherent compositionality of logical relations to know $\mathcal{C}[\llbracket t_1 \rrbracket] \approx \llbracket \mathcal{C}[\llbracket t_1 \rrbracket] \rrbracket$ and the same for t_2 . If the logical relations are constructed adequately, two terms are related if they equi-terminate. Thus, $\mathcal{C}[\llbracket t_1 \rrbracket] \Downarrow$ iff $\llbracket \mathcal{C}[\llbracket t_1 \rrbracket] \rrbracket \Downarrow$ and similarly for t_2 . In particular, this yields the implications (1) and (3) in Fig. 1. Since implication (2) follows directly from the hypothesis of contextual equivalence for $\llbracket t_1 \rrbracket$ and $\llbracket t_2 \rrbracket$, the proof for compiler correctness is finished.

Compiler security ($t_1 \approx_{ctx} t_2 \Rightarrow \llbracket t_1 \rrbracket \approx_{ctx} \llbracket t_2 \rrbracket$) requires that equivalent programs remain equivalent after compilation. This means that no matter what target-level manipulations are done on compiled programs, the programs must behave equivalently if the source programs were equivalent. This precludes all sorts of target-level attacks that break source-level guarantees.

If the source language is strong enough, it is possible to apply a strategy analogous to proving compiler correctness for proving compiler security, as depicted in Fig. 2. Given an arbitrary target-level context \mathcal{C} , we need to prove that $\mathcal{C}[\llbracket t_1 \rrbracket] \Downarrow$ implies $\mathcal{C}[\llbracket t_2 \rrbracket] \Downarrow$. In a sufficiently-powerful source language, we can construct a *back-translation* $\langle\langle \mathcal{C} \rangle\rangle$ for any target-level context \mathcal{C} that returns a valid source-level context \mathcal{C} . Using the same logical relation as above, it then suffices to prove that $\langle\langle \mathcal{C} \rangle\rangle \approx \mathcal{C}$ for any valid context. If we have that $\langle\langle \mathcal{C} \rangle\rangle \approx \mathcal{C}$ and $t_1 \approx \llbracket t_1 \rrbracket$, and similarly for t_2 , compositionality and adequacy of the logical relation yield implications (1) and (3) in the figure. The remaining implication (2) follows from the assumed contextual equivalence of t_1 and t_2 .

Constructing a back-translation of contexts is not easy, but it can be done if the source language is sufficiently expressive. Consider, for example, a compiler that translates terms from a simply-typed λ -calculus *with* recursive types ($\lambda^{\tau;\mu}$) to an untyped λ -calculus (λ^u). Constructing a back-translation of target-level contexts can be done based on a $\lambda^{\tau;\mu}$ type that can represent arbitrary λ^u values. Particularly, we can encode the untype of λ^u values in a type $UVal$ as follows:

$$UVal \stackrel{\text{def}}{=} \mu\alpha. \mathcal{B} \uplus (\alpha \times \alpha) \uplus (\alpha \uplus \alpha) \uplus (\alpha \rightarrow \alpha)$$

given that λ^u has base values of type \mathcal{B} , pairs, coproducts and lambdas. In other words, all λ^u values can be represented as $\lambda^{\tau;\mu}$ values of type $UVal$. We can then construct a back-translation of λ^u contexts to $\lambda^{\tau;\mu}$ contexts where the latter work with values in $UVal$ where the original λ^u context worked with arbitrary λ^u values.

Contributions of this paper If the types of the source language are not powerful enough to embed an encoding of target terms, is it possible to have a fully-abstract compiler between those languages?




Figure 3. Proving compiler security using an n -approximate back-translation. An $_$ subscript indicates *any* number of steps.

In this paper we answer positively to this question and develop a general proof technique for proving this. We instantiate this proof technique and develop a fully-abstract compiler from a simply-typed λ -calculus *without* recursive types (λ^{τ}) to an untyped λ -calculus (λ^u). With such a source language, we cannot construct a type like $UVal$ to represent the values that a λ^u context works with. Fortunately, we can solve this problem by observing that a fully accurate emulation is sufficient for the proof but in fact not necessary. An *approximate* back-translation is enough for the full-abstraction proof to work, without sacrificing the overall simplicity and elegance of the proof technique. The basic idea is depicted in Fig. 3. The differences from Fig. 2 are the use of asymmetric logical relations \lesssim and \gtrsim (also known as logical approximations) to express (roughly) that a term (or context) t terminates whenever t does ($t \gtrsim t$) and vice versa ($t \lesssim t$) and the addition of subscripts n where logical approximations hold only up to a limited number of steps n . Note that n in the Figure is defined as the number of steps in the evaluation $\mathcal{C}[\llbracket t_1 \rrbracket] \Downarrow_n$ and that we write $_$ for an unknown number of steps.

The proof starts, again, from an arbitrary target-level context \mathcal{C} and the knowledge that $\mathcal{C}[\llbracket t \rrbracket] \Downarrow_n$ and we call n the number of reduction steps in this execution. We then construct a λ^{τ} context $\langle\langle \mathcal{C} \rangle\rangle_n$ that satisfies two conditions. First, it approximates \mathcal{C} *up to* n steps: $\langle\langle \mathcal{C} \rangle\rangle_n \gtrsim_n \mathcal{C}$. This means that if $\mathcal{C}[\llbracket t \rrbracket]$ terminates in less than n steps then $\langle\langle \mathcal{C} \rangle\rangle_n[\llbracket t \rrbracket]$ will also terminate for a term t related to t . This, together with the knowledge that $t \gtrsim \llbracket t \rrbracket$, allows us to deduce implication (1) in the figure. As before, implication (2) follows directly from the contextual equivalence of t_1 and t_2 .

Then we use a second condition on the n -approximation $\langle\langle \mathcal{C} \rangle\rangle_n$, namely that it has to be *conservative*, to deduce implication (3). Intuitively, the source-level context produced by the n -approximation may diverge in situations where the original did not, but not vice versa. Intuitively, the divergence will occur when the precision n of approximate back-translation $\langle\langle \mathcal{C} \rangle\rangle_n$ is not sufficient for the context to accurately simulate the behavior of \mathcal{C} . This is expressed by the logical approximation $\langle\langle \mathcal{C} \rangle\rangle_n \lesssim \mathcal{C}$ which implies that if $\langle\langle \mathcal{C} \rangle\rangle_n[\llbracket t \rrbracket]$ terminates (in any number of steps), then so must $\mathcal{C}[\llbracket t \rrbracket]$. This allows us to deduce implication (3).

The advantage of this approximate back-translation approach is that it can be easier to construct a conservative approximate back-translation than a full one. For example, considering λ^{τ} without recursive types, we can construct a family of λ^{τ} types $UVal_n$, indexed by non-negative numbers n :

$$UVal_0 \stackrel{\text{def}}{=} U\text{unit}$$

$$UVal_{n+1} \stackrel{\text{def}}{=} U\text{unit} \uplus \mathcal{B} \uplus (UVal_n \times UVal_n) \uplus (UVal_n \uplus UVal_n) \uplus (UVal_n \rightarrow UVal_n).$$

Without giving full details here, $UVal_n$ is an n -level unfolding of $UVal$ with additional unit values at every level to represent failed approximations. This approximate version of $UVal$ is enough to construct a conservative n -approximate back-translation of an

untyped program context, and as such, it allows us to circumvent the lack of expressiveness of λ^τ without recursive types.

In order to make this approximate back-translation approach work, we need a way to formalise the relation between an untyped context and its approximate back-translation. However, it turns out that existing well-known techniques from the field of logical relations are almost directly applicable. Asymmetric logical relations (like $\langle\langle \mathcal{C} \rangle\rangle_n \lesssim \mathcal{C}$ above) are a well-established technique. More interestingly, the approximateness of the relation can very naturally be expressed using step-indexed logical relations. Despite this naturality, it appears that this use of step-indexing is novel. The technique is normally used as a way to construct well-founded logical relations and one is not actually interested in terms being related only up to a limited number of steps.

To summarise, the contributions of this work are:

- a new and general proof technique for proving compiler full-abstraction using asymmetric, cross-language logical relations and targeting untyped languages;
- an instantiation of that proof technique to fully-abstractly compile a simply-typed λ -calculus without recursive types to the untyped λ -calculus;
- a novel application of step-indexed logical relations for expressing approximateness of a back-translation.

This paper is structured as follows. Firstly it formalises the source and target languages λ^τ and λ^u (Section 2). Secondly it presents the cross-language logical relations that are used to express the relation between λ^τ terms and their compilations as well as between λ^u contexts and their back-translation (Section 3). We define the compiler in Section 4. It applies type erasure and dynamic type wrappers that enforce the requirements and guarantees of λ^τ types during execution. The paper then presents the approximate back-translation (Section 5) which is used when proving compiler full-abstraction (Section 6). Finally, we discuss the presented results (Section 7), related work (Section 8) and we conclude in Section 9.

2. Source and Target Languages

The source language λ^τ is presented in Fig. 4. It is a standard, strict, simply-typed λ -calculus with `Unit`, `Bool`, lambdas, product and sum types and a `fix` operator providing general recursion. The figure presents the syntax of terms t , values v , types τ , typing contexts Γ and evaluation contexts \mathcal{C} . Apart from the type and evaluation rules for `fix` $_{\tau_1 \rightarrow \tau_2}$, the typing rules and evaluation rules are standard. The evaluation rules use evaluation contexts to impose a strict evaluation order. The type and evaluation rule for `fix` $_{\tau_1 \rightarrow \tau_2}$ are somewhat special compared to a more standard definition (see e.g. [26]): the operator is restricted to function types and an additional η -expansion occurs during evaluation. This is because we have chosen to make `fix` model the Z fixed-point combinator (also known as the call-by-value Y combinator) [26, §5] rather than the Y combinator. The reason revolves around the compiler devised in this paper. The target language of that compiler is a *strict* untyped lambda calculus, where Y does not work but Z does and using Z in λ^u as well keeps the compiler simpler. Working with the more standard Y fixpoint combinator in λ^τ is probably possible but would require the compiler to use an encoding that would be pervasive but irrelevant to the subject of this paper.

The definition of λ^τ program contexts \mathcal{C} is omitted, they are a superset of evaluation contexts \mathcal{C} and they are essentially defined as λ^τ term that contain exactly one hole $[\]$ in place of a subterm. We also omit the typing judgement for program contexts $\vdash \mathcal{C} : \Gamma', \tau' \rightarrow \Gamma, \tau$, defined by inductive rules close to those for terms in Fig. 4. The judgement guarantees that substituting a well-typed

$$\begin{aligned}
\underline{W} &::= (k) \text{ with } k \in \mathbb{N} && \triangleright (k+1) \stackrel{\text{def}}{=} (k) \\
\text{lev}(k) &\stackrel{\text{def}}{=} \underline{W}.k && (k) \sqsupseteq (k') \stackrel{\text{def}}{=} k \leq k' \\
&&& \triangleright (0) \stackrel{\text{def}}{=} (0) && (k) \sqsupset_{\triangleright} (k') \stackrel{\text{def}}{=} k < k' \\
O(\underline{W})_{\lesssim} &\stackrel{\text{def}}{=} \left\{ (t, t) \mid \exists k \leq \text{lev}(\underline{W}), v, t \hookrightarrow^k v \Rightarrow \exists k', v, t \hookrightarrow^{k'} v \right\} \\
O(\underline{W})_{\gtrsim} &\stackrel{\text{def}}{=} \left\{ (t, t) \mid \exists k \leq \text{lev}(\underline{W}), v, t \hookrightarrow^k v \Rightarrow \exists k', v, t \hookrightarrow^{k'} v \right\}
\end{aligned}$$

Figure 6. Logical relations: Worlds.

term $\Gamma' \vdash t : \tau'$ in a well-typed context $\vdash \mathcal{C} : \Gamma', \tau' \rightarrow \Gamma, \tau$ produces a well-typed term $\Gamma \vdash \mathcal{C}[t] : \tau$.

Figure 5 presents the syntax, well-scopedness and evaluation rules for the target language λ^u : a standard untyped λ -calculus. The calculus has `unit`, booleans, lambdas, product and sum values, and produces an error term `wrong` in case of type errors (e.g. projecting from a non-pair value, case splitting on a non-sum value etc.). The well-scopedness rules are unsurprising and the evaluation rules again use evaluation contexts to impose a strict evaluation order. Again, we omit the definition of program contexts \mathcal{C} (expressions with a single hole in place of a subterm) and their well-scopedness judgement $\vdash \mathcal{C} : \Gamma' \rightarrow \Gamma$, whose inductive definition guarantees that substituting a well-scoped term $\Gamma' \vdash t$ for the hole produces a well-scoped result term $\Gamma \vdash \mathcal{C}[t]$.

The interested reader can find full formalisation and proofs in the companion technical report.

3. Logical relations

This section presents the Kripke, step-indexed logical relations that are used to prove compiler full-abstraction. Firstly, this section describes the specifications of the world used by the logical relation (Fig. 6). Then, it defines the logical relations (Fig. 7) and finally it proves standard properties that the relations enjoy. Note that part of these logical relations, namely the novel insights needed to provide compiler full-abstraction, are postponed until Section 5.2. The goal of this section is to provide an understanding of what it means for two terms to be related; this will be needed for understanding properties of the compiler in the following sections.

The cross-language logical relations used in this paper are roughly based on one by Hur and Dreyer [13]. Essentially, we instantiate their language-generic logical relations to λ^τ and λ^u and simplify the complexities deriving from the System F type system, public/private transitions, references and garbage collections.

Since we do not deal with mutable references, we use a very simple notion of worlds, consisting just of a step-index k that can be accessed with the `lev`(\cdot) function (Fig. 6). We define a \triangleright modality and a future world relation \sqsupseteq , expressing that future worlds allow less reduction steps to be taken. We define two different observation relations $O(\underline{W})_{\lesssim}$ and $O(\underline{W})_{\gtrsim}$. The former defines that a λ^τ term t approximates a λ^u term t if termination of the first in less than `lev`(\underline{W}) steps implies termination of the second (in an unknown number of steps). The latter requires the reverse. All of our logical relations will be defined in terms of either $O(\underline{W})_{\lesssim}$ or $O(\underline{W})_{\gtrsim}$. For definitions and lemmas or theorems that apply for both instantiations, we use the symbol \square as a metavariable that can be instantiated to either \lesssim and \gtrsim .

Figure 7 contains the definition of the logical relations. The first thing to note is that our logical relations are not indexed by λ^τ types τ , but by *pseudo-types* $\hat{\tau}$. The syntax for these pseudo-types contains all the constructs of λ^τ types, plus an additional kind of *token type* `EmuLDV $_{n,p}$` , indexed by a non-negative number n and a value $p ::= \text{precise} \mid \text{imprecise}$. This token type

$$\begin{aligned}
t &::= v \mid x \mid tt \mid t.1 \mid t.2 \mid \langle t, t \rangle \mid \text{inl } t \mid \text{inr } t \mid \text{case } t \text{ of inl } x_1 \mapsto t \mid \text{inr } x_2 \mapsto t \mid t; t \mid \text{if } t \text{ then } t \text{ else } t \mid \text{fix}_{\tau_1 \rightarrow \tau_2} t \\
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 \\
\tau &::= \text{Unit} \mid \text{Bool} \mid \tau \rightarrow \tau \mid \tau \times \tau \mid \tau \uplus \tau & \Gamma &::= \emptyset \mid \Gamma, x : \tau \\
C &::= [\] \mid C t \mid v C \mid C.1 \mid C.2 \mid \langle C, t \rangle \mid \langle v, C \rangle \mid \text{inl } C \mid \text{inr } C \mid \text{case } C \text{ of inl } x_1 \mapsto t_1 \mid \text{inr } x_2 \mapsto t_2 \mid C; t \mid \text{if } C \text{ then } t \text{ else } t \\
&\quad \mid \text{fix}_{\tau_1 \rightarrow \tau_2} C
\end{aligned}$$

$$\begin{array}{c}
\frac{}{\Gamma \vdash \text{unit} : \text{Unit}} \quad \frac{}{\Gamma \vdash \text{true} : \text{Bool}} \quad \frac{(x : \tau) \in \Gamma}{\Gamma \vdash x : \tau} \quad \frac{\Gamma, (x : \tau) \vdash t : \tau'}{\Gamma \vdash \lambda x : \tau. t : \tau'} \quad \frac{\Gamma \vdash t_1 : \tau_1 \quad \Gamma \vdash t_2 : \tau_2}{\Gamma \vdash \langle t_1, t_2 \rangle : \tau_1 \times \tau_2} \\
\frac{\Gamma \vdash t : \tau_1 \times \tau_2}{\Gamma \vdash t.1 : \tau_1} \quad \frac{\Gamma \vdash t : \tau' \rightarrow \tau \quad \Gamma \vdash t' : \tau'}{\Gamma \vdash t t' : \tau} \quad \frac{\Gamma \vdash t : \tau}{\Gamma \vdash \text{inl } t : \tau \uplus \tau'} \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} \\
\frac{\Gamma \vdash t : \tau_1 \uplus \tau_2 \quad \Gamma, (x_1 : \tau_1) \vdash t_1 : \tau \quad \Gamma, (x_2 : \tau_2) \vdash t_2 : \tau}{\Gamma \vdash \text{case } t \text{ of inl } x_1 \mapsto t_1 \mid \text{inr } x_2 \mapsto t_2 : \tau} \quad \frac{\Gamma \vdash t : \text{Bool} \quad \Gamma \vdash t_1 : \tau \quad \Gamma \vdash t_2 : \tau}{\Gamma \vdash \text{if } t \text{ then } t_1 \text{ else } t_2 : \tau} \quad \frac{\Gamma, \vdash t_1 : \text{Unit} \quad \Gamma, \vdash t_2 : \tau}{\Gamma \vdash t_1; t_2 : \tau} \\
\frac{t \hookrightarrow t'}{C[t] \hookrightarrow C[t']} \quad \frac{}{(\lambda x : \tau. t) v \hookrightarrow t[v/x]} \quad \frac{}{\langle v_1, v_2 \rangle.1 \hookrightarrow v_1} \quad \frac{\text{case inl } v \text{ of inl } x_1 \mapsto t_1 \mid \text{inr } x_2 \mapsto t_2 \hookrightarrow t_1[v/x_1]}{} \\
\frac{v \equiv \text{true} \Rightarrow t' \equiv t_1 \quad v \equiv \text{false} \Rightarrow t' \equiv t_2}{\text{if } v \text{ then } t_1 \text{ else } t_2 \hookrightarrow t'} \quad \frac{}{\text{unit}; t \hookrightarrow t} \quad \frac{\text{fix}_{\tau_1 \rightarrow \tau_2} (\lambda x : \tau_1 \rightarrow \tau_2. t) \hookrightarrow t[(\lambda y : \tau_1. \text{fix}_{\tau_1 \rightarrow \tau_2} (\lambda x : \tau_1 \rightarrow \tau_2. t) y)/x]}{}
\end{array}$$

Figure 4. Syntax, static and dynamic semantics of the source language λ^τ (selection of).

$$\begin{aligned}
t &::= \text{unit} \mid \text{true} \mid \text{false} \mid \lambda x. t \mid x \mid tt \mid t.1 \mid t.2 \mid \langle t, t \rangle \mid \text{inl } t \mid \text{inr } t \mid \text{case } t \text{ of inl } x \mapsto t \mid \text{inr } x \mapsto t \mid t; t \mid \text{if } t \text{ then } t \text{ else } t \mid \text{wrong} \\
v &::= \text{unit} \mid \text{true} \mid \text{false} \mid \lambda x. t \mid \langle v, v \rangle \mid \text{inl } v \mid \text{inr } v & \Gamma &::= \emptyset \mid \Gamma, x \\
C &::= [\] \mid C t \mid v C \mid C.1 \mid C.2 \mid \langle C, t \rangle \mid \langle v, C \rangle \mid \text{inl } C \mid \text{inr } C \mid \text{case } C \text{ of inl } x_1 \mapsto t_1 \mid \text{inr } x_2 \mapsto t_2 \mid C; t \mid \text{if } C \text{ then } t \text{ else } t
\end{aligned}$$

$$\begin{array}{c}
\frac{t \hookrightarrow t'}{C[t] \hookrightarrow C[t']} \quad \frac{}{C[\text{wrong}] \hookrightarrow \text{wrong}} \quad \frac{}{(\lambda x. t) v \hookrightarrow t[v/x]} \quad \frac{}{\langle v_1, v_2 \rangle.1 \hookrightarrow v_1} \\
\frac{\text{case inl } v \text{ of inl } x_1 \mapsto t_1 \mid \text{inr } x_2 \mapsto t_2 \hookrightarrow t_1[v/x_1]}{} \quad \frac{v \equiv \text{true} \Rightarrow t' \equiv t_1 \quad v \equiv \text{false} \Rightarrow t' \equiv t_2 \quad (v \neq \text{true} \wedge v \neq \text{false}) \Rightarrow t' \equiv \text{wrong}}{\text{if } v \text{ then } t_1 \text{ else } t_2 \hookrightarrow t'} \quad \frac{}{\text{unit}; t \hookrightarrow t}
\end{array}$$

Figure 5. Syntax and dynamic semantics of the target language λ^u (selection of).

is not an λ^τ type; it is needed because of the approximate back-translation. When necessary, we use a function `repEmul()` for converting a pseudo-type to a λ^τ type. The function replaces all occurrences of $\text{EmulDV}_{n,p}$ with a concrete λ^τ type. We postpone the definitions and explanations of $\text{EmulDV}_{n,p}$ and of $\mathcal{V}[\llbracket \text{EmulDV}_{n,p} \rrbracket]_\square$ to Section 5.2, after we have given some more information about the back-translation. We will sometimes silently use a normal type where a pseudo-type is expected, which makes sense since the syntax for the latter is a superset of the former.

The value relation $\mathcal{V}[\hat{\tau}]_\square$ is defined by induction on the pseudo-type. Most definitions are quite standard. All cases require related terms to be in the `of` type relation, which requires well-typedness of the λ^τ term and an appropriate shape for the λ^u value. `Unit` and `Bool` values are related in any world iff they are the same base value. Pair values are related if both are pairs and the corresponding components are related in strictly future worlds at the appropriate pseudo-type. Similarly, sum values are related if they are both of either the form `inl ...` or `inr ...` and if the contained values are related in strictly future worlds at the appropriate pseudo-type.

Finally, function values are related if they have the right type, if both are lambdas and if substituting related values in the body yields related terms in any strictly future world.

The relation on values, evaluation contexts and terms are defined mutually recursively, using a technique known as biorthogonality (see, e.g., [6]). So, evaluation contexts are related in a world if plugging in related values in any future world yields related observations. Similarly, terms are related if plugging the terms in related evaluation contexts yields related observations. Relation $\mathcal{G}[\Gamma]_\square$ relates substitutions instantiating a context Γ , which simply requires that substitutions for all variables in the context are related at their types. For open terms, we define a logical relation $\hat{\Gamma} \vdash t \square_n t : \hat{\tau}$. This relation expresses that an open λ^τ term t is related up to n steps to an open λ^u term t at pseudo-type $\hat{\tau}$ in pseudo-context $\hat{\Gamma}$ if the first is well-typed and if closing t and t with substitutions related at pseudo-context $\hat{\Gamma}$ produces terms related at pseudo-type $\hat{\tau}$, in any world \underline{W} such that $\text{lev}(\underline{W}) \leq n$. If $\hat{\Gamma} \vdash t \square_n t : \hat{\tau}$ for any n , then we write

Pseudo-types $\hat{\tau}$, pseudo-contexts $\hat{\Gamma}$, $\text{ofType}(\cdot)$ and $\text{repEmul}(\cdot)$.

$$\begin{aligned} \hat{\tau} &::= \text{Bool} \mid \text{Unit} \mid \hat{\tau} \times \hat{\tau} \mid \hat{\tau} \uplus \hat{\tau} \mid \hat{\tau} \rightarrow \hat{\tau} \mid \text{EmulDV}_{n,p} & \hat{\Gamma} &::= \emptyset \mid \hat{\Gamma}, \mathbf{x} : \hat{\tau} \\ \text{repEmul}(\hat{\tau}) &\stackrel{\text{def}}{=} \dots \text{ (to be defined later, in Fig. 12)} \\ \text{ofType}(\hat{\tau}) &\stackrel{\text{def}}{=} \{ \mathbf{v} \mid \emptyset \vdash \mathbf{v} : \text{repEmul}(\hat{\tau}) \} \\ \text{ofType}(\hat{\tau}) &\stackrel{\text{def}}{=} \left\{ \mathbf{v} \left| \begin{array}{l} \mathbf{v} = \text{unit} \\ \mathbf{v} = \text{true} \text{ or } \mathbf{v} = \text{false} \\ \exists t. \mathbf{v} = \lambda x. t \\ \exists v_1 \in \text{ofType}(\hat{\tau}_1), v_2 \in \text{ofType}(\hat{\tau}_2). \mathbf{v} = \langle v_1, v_2 \rangle \\ \exists v_1 \in \text{ofType}(\hat{\tau}_1). \mathbf{v} = \text{inl } v_1 \text{ or } \exists v_2 \in \text{ofType}(\hat{\tau}_2). \mathbf{v} = \text{inr } v_2 \end{array} \right. \begin{array}{l} \text{if } \hat{\tau} = \text{Unit} \\ \text{if } \hat{\tau} = \text{Bool} \\ \text{if } \exists \hat{\tau}_1, \hat{\tau}_2. \hat{\tau} = \hat{\tau}_1 \rightarrow \hat{\tau}_2 \\ \text{if } \exists \hat{\tau}_1, \hat{\tau}_2. \hat{\tau} = \hat{\tau}_1 \times \hat{\tau}_2 \\ \text{if } \exists \hat{\tau}_1, \hat{\tau}_2. \hat{\tau} = \hat{\tau}_1 \uplus \hat{\tau}_2 \end{array} \right\} \\ \text{ofType}(\hat{\tau}) &\stackrel{\text{def}}{=} \{ (\mathbf{v}, \mathbf{v}) \mid \mathbf{v} \in \text{ofType}(\hat{\tau}) \wedge \mathbf{v} \in \text{ofType}(\hat{\tau}) \} \end{aligned}$$

Logical relations for values ($\mathcal{V}[\cdot]_{\square}$), contexts ($\mathcal{K}[\cdot]_{\square}$), terms ($\mathcal{E}[\cdot]_{\square}$) and environments ($\mathcal{G}[\cdot]_{\square}$).

$$\begin{aligned} \triangleright R &\stackrel{\text{def}}{=} \{ (\underline{W}, \mathbf{v}, \mathbf{v}) \mid \text{lev}(\underline{W}) > 0 \Rightarrow (\triangleright \underline{W}, \mathbf{v}, \mathbf{v}) \in R \} \\ \mathcal{V}[\text{Unit}]_{\square} &\stackrel{\text{def}}{=} \{ (\underline{W}, \mathbf{v}, \mathbf{v}) \mid \mathbf{v} = \text{unit} \text{ and } \mathbf{v} = \text{unit} \} \\ \mathcal{V}[\text{Bool}]_{\square} &\stackrel{\text{def}}{=} \{ (\underline{W}, \mathbf{v}, \mathbf{v}) \mid \exists v \in \{\text{true}, \text{false}\}. \mathbf{v} = v \text{ and } \mathbf{v} = v \} \\ \mathcal{V}[\hat{\tau}' \rightarrow \hat{\tau}]_{\square} &\stackrel{\text{def}}{=} \left\{ (\underline{W}, \mathbf{v}, \mathbf{v}) \left| \begin{array}{l} (\mathbf{v}, \mathbf{v}) \in \text{ofType}(\hat{\tau}' \rightarrow \hat{\tau}) \text{ and } \exists t, \mathbf{t}. \mathbf{v} = \lambda x. \text{repEmul}(\hat{\tau}'). \mathbf{t} \text{ and } \mathbf{v} = \lambda x. \mathbf{t} \text{ and} \\ \forall \underline{W}' \triangleright \underline{W}, (\underline{W}', \mathbf{v}', \mathbf{v}') \in \mathcal{V}[\hat{\tau}']_{\square}. (\underline{W}', \mathbf{t}[\mathbf{v}'/\mathbf{x}], \mathbf{t}[\mathbf{v}'/\mathbf{x}]) \in \mathcal{E}[\hat{\tau}]_{\square} \end{array} \right. \right\} \\ \mathcal{V}[\hat{\tau}_1 \times \hat{\tau}_2]_{\square} &\stackrel{\text{def}}{=} \left\{ (\underline{W}, \mathbf{v}, \mathbf{v}) \left| \begin{array}{l} (\mathbf{v}, \mathbf{v}) \in \text{ofType}(\hat{\tau}_1 \times \hat{\tau}_2) \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 \text{ and } \mathbf{v} = \langle \mathbf{v}_1, \mathbf{v}_2 \rangle \text{ and} \\ (\underline{W}, \mathbf{v}_1, \mathbf{v}_1) \in \triangleright \mathcal{V}[\hat{\tau}_1]_{\square} \text{ and } (\underline{W}, \mathbf{v}_2, \mathbf{v}_2) \in \triangleright \mathcal{V}[\hat{\tau}_2]_{\square} \end{array} \right. \right\} \\ \mathcal{V}[\hat{\tau}_1 \uplus \hat{\tau}_2]_{\square} &\stackrel{\text{def}}{=} \left\{ (\underline{W}, \mathbf{v}, \mathbf{v}) \left| \begin{array}{l} (\mathbf{v}, \mathbf{v}) \in \text{ofType}(\hat{\tau}_1 \uplus \hat{\tau}_2) \text{ and} \\ \exists \mathbf{v}', \mathbf{v}'. (\underline{W}, \mathbf{v}', \mathbf{v}') \in \triangleright \mathcal{V}[\hat{\tau}_1]_{\square} \text{ and } \mathbf{v} = \text{inl } \mathbf{v}' \text{ and } \mathbf{v} = \text{inl } \mathbf{v}' \text{ or} \\ \exists \mathbf{v}', \mathbf{v}'. (\underline{W}, \mathbf{v}', \mathbf{v}') \in \triangleright \mathcal{V}[\hat{\tau}_2]_{\square} \text{ and } \mathbf{v} = \text{inr } \mathbf{v}' \text{ and } \mathbf{v} = \text{inr } \mathbf{v}' \end{array} \right. \right\} \end{aligned}$$

$\mathcal{V}[\text{EmulDV}_{n,p}]_{\square} \stackrel{\text{def}}{=} \dots$ (to be defined later, in Fig. 11)

$$\begin{aligned} \mathcal{K}[\hat{\tau}]_{\square} &\stackrel{\text{def}}{=} \{ (\underline{W}, \underline{C}, \underline{C}) \mid \forall \underline{W}' \triangleright \underline{W}, (\underline{W}', \mathbf{v}, \mathbf{v}) \in \mathcal{V}[\hat{\tau}]_{\square}. (\underline{C}[\mathbf{v}], \underline{C}[\mathbf{v}]) \in \mathcal{O}(\underline{W}')_{\square} \} \\ \mathcal{E}[\hat{\tau}]_{\square} &\stackrel{\text{def}}{=} \{ (\underline{W}, \underline{t}, \underline{t}) \mid \forall (\underline{W}, \underline{C}, \underline{C}) \in \mathcal{K}[\hat{\tau}]_{\square}. (\underline{C}[\underline{t}], \underline{C}[\underline{t}]) \in \mathcal{O}(\underline{W})_{\square} \} \end{aligned}$$

Logical relations for substitutions, open terms and program contexts.

$$\begin{aligned} \mathcal{G}[\emptyset]_{\square} &\stackrel{\text{def}}{=} \{ (\underline{W}, \emptyset, \emptyset) \} \\ \mathcal{G}[\hat{\Gamma}, (\mathbf{x} : \hat{\tau})]_{\square} &\stackrel{\text{def}}{=} \{ (\underline{W}, \gamma[\mathbf{x} \mapsto \mathbf{v}], \gamma[\mathbf{x} \mapsto \mathbf{v}]) \mid (\underline{W}, \gamma, \gamma) \in \mathcal{G}[\hat{\Gamma}]_{\square} \text{ and } (\underline{W}, \mathbf{v}, \mathbf{v}) \in \mathcal{V}[\hat{\tau}]_{\square} \} \\ \hat{\Gamma} \vdash \mathbf{t} \square_n \mathbf{t} : \hat{\tau} &\stackrel{\text{def}}{=} \text{repEmul}(\hat{\Gamma}) \vdash \mathbf{t} : \text{repEmul}(\hat{\tau}) \text{ and } \forall \underline{W}. \text{lev}(\underline{W}) \leq n \Rightarrow \forall (\underline{W}, \gamma, \gamma) \in \mathcal{G}[\hat{\Gamma}]_{\square}. (\underline{W}, \mathbf{t}\gamma, \mathbf{t}\gamma) \in \mathcal{E}[\hat{\tau}]_{\square} \\ \hat{\Gamma} \vdash \mathbf{t} \square \mathbf{t} : \hat{\tau} &\stackrel{\text{def}}{=} \hat{\Gamma} \vdash \mathbf{t} \square_n \mathbf{t} : \hat{\tau} \text{ for all } n \\ \vdash \underline{c} \square \underline{c} : \hat{\Gamma}', \hat{\tau}' \rightarrow \hat{\Gamma}, \hat{\tau} &\stackrel{\text{def}}{=} \vdash \underline{c} : \text{repEmul}(\hat{\Gamma}'), \text{repEmul}(\hat{\tau}') \rightarrow \text{repEmul}(\hat{\Gamma}), \text{repEmul}(\hat{\tau}) \text{ and} \\ &\quad \text{for all } \mathbf{t}, \mathbf{t}. \text{ if } \hat{\Gamma}' \vdash \mathbf{t} \square \mathbf{t} : \hat{\tau}', \text{ then } \hat{\Gamma} \vdash \underline{c}[\mathbf{t}] \square \underline{c}[\mathbf{t}] : \hat{\tau} \end{aligned}$$

Figure 7. Logical relations (partial, the missing definition can be found in Figs. 11 and 12).

$\hat{\Gamma} \vdash \mathbf{t} \square \mathbf{t} : \hat{\tau}$. Finally, we define a logical relation for program contexts $\vdash \underline{c} \square \underline{c} : \hat{\Gamma}', \hat{\tau}' \rightarrow \hat{\Gamma}, \hat{\tau}$ which requires that substituting terms related at the appropriate pseudo-type type produces terms related at the appropriate pseudo-type.

It is interesting to note that the simple type system of our source calculus does not actually present a technical need for the use of step-indexing. Because there are no recursive types or general references, it is a simple enough system that we can give well-founded logical relations without any step-indexing. However, as mentioned before, we use step-indexing for a different reason than other work: not for constructing a well-founded logical relation, but for stating that two terms are related *only up to a certain number of steps*. More details follow in Section 5.

These logical relations are constructed so that termination of one implies termination of the other, according to the direction of the approximation (\lesssim or \gtrsim , Lemma 1).

Lemma 1 (Adequacy for \lesssim and \gtrsim). *If $\emptyset \vdash \mathbf{t} \lesssim_n \mathbf{t} : \tau$, and if $\mathbf{t} \leftrightarrow^m \mathbf{v}$ with $n \geq m$, then also $\mathbf{t} \Downarrow$. If $\emptyset \vdash \mathbf{t} \gtrsim_n \mathbf{t} : \tau$ and if $\mathbf{t} \leftrightarrow^m \mathbf{v}$ with $n \geq m$, then also $\mathbf{t} \Downarrow$.*

4. The Compiler

This section presents our compiler from λ^{τ} to λ^u . The compiler proceeds in two passes: type erasure (Fig. 8) and dynamic type-checking wrappers (Fig. 9).

$$\begin{aligned}
\text{fix} &\stackrel{\text{def}}{=} \lambda f. (\lambda x. f (\lambda y. x \times y)) (\lambda x. f (\lambda y. x \times y)) \\
\text{erase}(\text{unit}) &\stackrel{\text{def}}{=} \text{unit} \quad \text{erase}((t_1, t_2)) \stackrel{\text{def}}{=} (\text{erase}(t_1), \text{erase}(t_2)) \\
\text{erase}(\text{false}) &\stackrel{\text{def}}{=} \text{false} \quad \text{erase}(t_1; t_2) \stackrel{\text{def}}{=} \text{erase}(t_1); \text{erase}(t_2) \\
\text{erase}(x) &\stackrel{\text{def}}{=} x \quad \text{erase}(t_1 t_2) \stackrel{\text{def}}{=} \text{erase}(t_1) \text{erase}(t_2) \\
\text{erase}(t.1) &\stackrel{\text{def}}{=} \text{erase}(t).1 \\
\text{erase}(t.2) &\stackrel{\text{def}}{=} \text{erase}(t).2 \quad \text{erase}(\text{true}) \stackrel{\text{def}}{=} \text{true} \\
\text{erase}(\lambda x : \tau. t) &\stackrel{\text{def}}{=} \lambda x. \text{erase}(t) \\
\text{erase}(\text{inl } t) &\stackrel{\text{def}}{=} \text{inl } \text{erase}(t) \quad \text{erase}(\text{fix}_{\tau_1 \rightarrow \tau_2} t) \stackrel{\text{def}}{=} \text{fix } \text{erase}(t) \\
\text{erase}(\text{inr } t) &\stackrel{\text{def}}{=} \text{inr } \text{erase}(t) \\
\text{erase}(\text{if } t \text{ then } t_1 \text{ else } t_2) &\stackrel{\text{def}}{=} \text{if } \text{erase}(t) \text{ then } \text{erase}(t_1) \\
&\quad \text{else } \text{erase}(t_2) \\
\text{erase}(\text{case } t \text{ of inl } x_1 \mapsto t_1 \mid \text{inr } x_2 \mapsto t_2) &\stackrel{\text{def}}{=} \\
&\quad \text{case } \text{erase}(t) \text{ of } \left\{ \begin{array}{l} \text{inl } x_1 \mapsto \text{erase}(t_1) \\ \text{inr } x_2 \mapsto \text{erase}(t_2) \end{array} \right.
\end{aligned}$$

Figure 8. Type erasure: the first pass of the compiler.

The erasure function is called `erase`; it converts all λ^τ constructs to the corresponding λ^u constructs. $\text{fix}_{\tau_1 \rightarrow \tau_2}$ is erased to a λ^u definition of the Z combinator `fix`.

The `erase` function can be considered as a compiler, but it is only a correct compiler, not a fully-abstract one, as explained in Example 4.1.

Example 4.1 (Erasure is correct but not secure [11, 24]). Consider the following, contextually equivalent λ^τ functions of type $\text{Unit} \rightarrow \text{Unit}$:

$$\lambda x : \text{Unit}. x \quad \simeq_{\text{ctx}} \quad \lambda x : \text{Unit}. \text{unit}$$

The `erase` function will map these to the following λ^u functions:

$$\lambda x. x \quad \not\simeq_{\text{ctx}} \quad \lambda x. \text{unit}$$

The results of `erase` are *not* contextually equivalent, essentially because applying them to a non-unit value like `true` will produce `true` for the left lambda and `unit` for the right lambda. In this example, contextual equivalence is not preserved because the original functions are only defined for `Unit` values, but their compilations can be applied to other values too.

Lemma 2 states that every λ^τ term is related to its erased term at its type.

Lemma 2 (Erasure is semantics-preserving (for terms)). *If $\Gamma \vdash t : \tau$, then $\Gamma \vdash t \sqsubseteq \text{erase}(t) : \tau$.*

An analogous result applies to program contexts:

Lemma 3 (Erasure is semantics-preserving (for context)). *For all \mathcal{C} , if $\vdash \mathcal{C} : \Gamma', \tau' \rightarrow \Gamma, \tau$ then $\vdash \mathcal{C} \sqsubseteq \text{erase}(\mathcal{C}) : \Gamma', \tau' \rightarrow \Gamma, \tau$.*

One should intuitively understand this result as “`t` behaves the same as `erase(t)` when both are treated as values of type τ ”. The result does not specify what happens when we treat `t` as a value of a different type, like we did in Example 4.1 to demonstrate a full abstraction failure. Intuitively, it only specifies a kind of *correctness* of the `erase` function, not its *security*.

Remember that a fully-abstract compiler must protect terms from being used in ways that are not allowed by their type, as in Example 4.1. This is taken care of by the second pass of the compiler.

We construct a family of dynamic typechecking wrappers `protect $_\tau$` and `confine $_\tau$` . `protect $_\tau$` is a λ^u lambda term that wraps

$$\begin{aligned}
\text{protect}_{\text{Unit}} &\stackrel{\text{def}}{=} \lambda x. x & \text{protect}_{\text{Bool}} &\stackrel{\text{def}}{=} \lambda x. x \\
\text{protect}_{\tau_1 \times \tau_2} &\stackrel{\text{def}}{=} \lambda y. (\text{protect}_{\tau_1} y.1, \text{protect}_{\tau_2} y.2) \\
\text{protect}_{\tau_1 \sqcup \tau_2} &\stackrel{\text{def}}{=} \lambda y. \text{case } y \text{ of } \left\{ \begin{array}{l} \text{inl } x \mapsto \text{inl } (\text{protect}_{\tau_1} x) \\ \text{inr } x \mapsto \text{inr } (\text{protect}_{\tau_2} x) \end{array} \right. \\
\text{protect}_{\tau_1 \rightarrow \tau_2} &\stackrel{\text{def}}{=} \lambda y. \lambda x. \text{protect}_{\tau_2} (y (\text{confine}_{\tau_1} x)) \\
\text{confine}_{\text{Unit}} &\stackrel{\text{def}}{=} \lambda y. (y; \text{unit}) \\
\text{confine}_{\text{Bool}} &\stackrel{\text{def}}{=} \lambda y. \text{if } y \text{ then } \text{true} \text{ else } \text{false} \\
\text{confine}_{\tau_1 \times \tau_2} &\stackrel{\text{def}}{=} \lambda y. (\text{confine}_{\tau_1} y.1, \text{confine}_{\tau_2} y.2) \\
\text{confine}_{\tau_1 \sqcup \tau_2} &\stackrel{\text{def}}{=} \lambda y. \text{case } y \text{ of } \left\{ \begin{array}{l} \text{inl } x \mapsto \text{inl } (\text{confine}_{\tau_1} x) \\ \text{inr } x \mapsto \text{inr } (\text{confine}_{\tau_2} x) \end{array} \right. \\
\text{confine}_{\tau_1 \rightarrow \tau_2} &\stackrel{\text{def}}{=} \lambda y. \lambda x. \text{confine}_{\tau_2} (y (\text{protect}_{\tau_1} x))
\end{aligned}$$

Figure 9. Dynamic type checking wrappers: the second pass of the compiler.

an argument to enforce that it can only *be used* in ways that are valid according to type τ , as often done in secure compilation work [4, 8, 11, 24]. Dually, `confine $_\tau$` wraps its argument so that it can only *behave* in ways that are valid according to type τ . In the definition, the cases for product and coproduct types simply recursively descend on their subterms preserving the expected syntax of a product or coproduct argument. Protecting at a function type means wrapping the function to confine its arguments and protect its results, and dually for confining at a function type. Finally, protecting at a base type (i.e., `Unit` or `Bool`) does nothing, simply because there is nothing one can do to a base value that is not allowed by its type. Confining a value at a base type is more interesting. Both for `Unit` and `Bool` values, we use the value in such a way that will only work when the value is actually of the correct type. If it is, we return the original value, otherwise the term will reduce to `wrong`.

Example 4.2 (Protect and confine make a term secure). Consider the protect wrapper for type $\text{Unit} \rightarrow \text{Unit}$ `protect $_{\text{Unit} \rightarrow \text{Unit}}$` , which is (roughly) equal to $\lambda y. \lambda x. y (x; \text{unit})$. Applying that wrapper to a function `f` (i.e. `protect $_{\text{Unit} \rightarrow \text{Unit}}$ f`) reduces to $\lambda x. f (x; \text{unit})$. Applying this value to a non-`unit` value will simply evaluate to `wrong`, therefore addressing the issues of Example 4.1.

For the second pass of the compiler, Lemma 4 holds.

Lemma 4 (Protect and confine are semantics-preserving). *If $\Gamma \vdash t \sqsubseteq_n t : \tau$ then $\Gamma \vdash t \sqsubseteq_n \text{protect}_\tau t : \tau$ and $\Gamma \vdash t \sqsubseteq_n \text{confine}_\tau t : \tau$.*

Lemma 4 states that if `t` is related to `t` at type τ , then adding a `protect $_\tau$` or `confine $_\tau$` wrapper around `t` does not change that. In other words, the wrappers do not change the behaviour of `t` as long as they are treated as values of type τ . In Section 5.5, we will have more to say about the security of the wrappers.

This section concludes with the definition of the compiler used in this paper.

Definition 1 (The $\llbracket \cdot \rrbracket$ compiler). *If $\Gamma \vdash t : \tau$, then `t` is compiled to $\llbracket t \rrbracket$ and: $\llbracket t \rrbracket \stackrel{\text{def}}{=} \text{protect}_\tau (\text{erase}(t))$.*

Lemmas 2 and 4 about the first and second pass of the compiler can be combined into Lemma 5 to obtain that a λ^τ term of type τ behaves like its compilation when both are treated as terms of type τ .

Lemma 5 ($\llbracket \cdot \rrbracket$ is semantics-preserving). *For all t , if $\Gamma \vdash t : \tau$ then $\Gamma \vdash t \square \llbracket t \rrbracket : \tau$.*

5. Approximate Back-Translation

This section presents the core idea of our proof technique: the approximate back-translation. As explained in Section 1, the idea is to translate a target language program context \mathcal{C} to a source language program context $\llbracket \mathcal{C} \rrbracket_n$ which conservatively n -approximates \mathcal{C} . Intuitively, this means that $\llbracket \mathcal{C} \rrbracket_n$ behaves like \mathcal{C} for up to n steps but it may diverge in cases where the original did not if \mathcal{C} takes more than n steps. We will make this more precise in Section 5.2.

At the core of the approximate back-translation is the λ^τ type $UVal_n$. The type is essentially a λ^τ encoding of the untype of λ^u . Where the untyped context \mathcal{C} manipulates arbitrary λ^u values, its emulation $\llbracket \mathcal{C} \rrbracket_n$ manipulates values of type $UVal_n$. Section 5.1 defines $UVal_n$ and the basic tools (constructors and destructors) for working with it. To explain how values in $UVal_n$ model values in λ^u , Section 5.2 fills in the missing piece of the logical relations of Fig. 7 by defining $\mathcal{V}[\llbracket \text{EmuLDV}_{n;p} \rrbracket]_\square$.

The type $UVal_n$ is sufficiently large to contain n -approximations of λ^u values. However, it also contains approximations of λ^u values up to less than n steps. This is crucial, as for a term to be well-typed the accuracy of the approximation can be less than n . In these cases values of type $UVal_n$ will be *downgraded* to a type $UVal_m$ with $m < n$. Dually, there will be cases where some values need to *upgrade*. Section 5.3 defines functions to perform value upgrading and downgrading.

With the definition of upgrading and downgrading, we have defined all the machinery that revolves around $UVal_n$. Section 5.4 constructs the function responsible for the emulating a context, emulate_n , such that it translates a λ^u term t into a λ^τ term of type $UVal_n$. This function is easily extended to work with program contexts, producing contexts with hole of type $UVal_n$ as expected.

However, remember from Fig. 3 in Section 1 that the goal of the back-translation is generating a context $\llbracket \mathcal{C} \rrbracket_n$ whose hole can be filled with λ^τ terms t_1 and t_2 . However, the type of t_1 and t_2 is not $UVal_n$ but an arbitrary λ^τ type τ . Thus, there is a type mismatch between the hole of the emulated context $\text{emulate}_n(\mathcal{C})$ and the terms that we want to plug in there. Since the emulated contexts work with $UVal_n$ values, we need a function that wraps terms of an arbitrary type τ into a value of type $UVal_n$. This is precisely what Section 5.5 defines, namely a function $\text{inject}_{\tau;n}$ of type $\tau \rightarrow UVal_n$.

Finally, Section 5.6 defines the approximate back-translation function $\langle \cdot \rangle_{\tau;n}$, mapping a λ^u context \mathcal{C} to a λ^τ context $\llbracket \mathcal{C} \rrbracket_{\tau;n}$. The additional index τ w.r.t. earlier discussions is needed to introduce an appropriate call to $\text{inject}_{\tau;n}$ as discussed above, so that the hole of $\llbracket \mathcal{C} \rrbracket_{\tau;n}$ is of type τ . Plugging a term t_1 in $\llbracket \mathcal{C} \rrbracket_{\tau;n}$ n -approximates plugging in the compilation $\llbracket t_1 \rrbracket$ in context \mathcal{C} .

Right after the definition of each of the concepts discussed above (*downgrade*, *upgrade*, $\text{inject}_{\tau;n}$ and emulate_n), this section formalises the results about their behaviour. These results are expressed in terms of the logical relations of Fig. 7 and of the $\text{EmuLDV}_{n;p}$ pseudo-type; they will be used to prove compiler security in Section 6.

5.1 $UVal$ and its Tools

The family of types $UVal_n$ is defined as follows:

$$UVal_0 \stackrel{\text{def}}{=} \text{Unit}$$

$$UVal_{n+1} \stackrel{\text{def}}{=} \text{Unit} \uplus \text{Unit} \uplus \text{Bool} \uplus (UVal_n \times UVal_n) \uplus (UVal_n \uplus UVal_n) \uplus (UVal_n \rightarrow UVal_n)$$

$UVal_n$ is the type that emulated λ^u terms have when back-translated into λ^τ . For every n , $UVal_n$ is clearly a valid λ^τ type. At

$$\begin{aligned} \text{in}_{\text{unk},n} &: UVal_{n+1} \\ \text{in}_{\text{unit},n} &: \text{Unit} \rightarrow UVal_{n+1} & \text{unk}_n &: UVal_n \\ \text{in}_{\text{bool},n} &: \text{Bool} \rightarrow UVal_{n+1} & \text{unk}_0 &\stackrel{\text{def}}{=} \text{Unit} \\ \text{in}_{\times,n} &: (UVal_n \times UVal_n) \rightarrow UVal_{n+1} & \text{unk}_{n+1} &\stackrel{\text{def}}{=} \text{in}_{\text{unk},n} \\ \text{in}_{\uplus,n} &: (UVal_n \uplus UVal_n) \rightarrow UVal_{n+1} \\ \text{in}_{\rightarrow,n} &: (UVal_n \rightarrow UVal_n) \rightarrow UVal_{n+1} \\ \text{omega}_\tau &: \tau & \text{omega}_\tau &\stackrel{\text{def}}{=} \text{fix}_{\text{Unit} \rightarrow \tau} (\lambda x : \text{Unit} \rightarrow \tau. x) \text{unit} \\ \text{case}_{\text{unit},n} &: UVal_{n+1} \rightarrow \text{Unit} & \text{case}_{\text{bool},n} &: UVal_{n+1} \rightarrow \text{Bool} \\ \text{case}_{\times,n} &: UVal_{n+1} \rightarrow (UVal_n \times UVal_n) \\ \text{case}_{\uplus,n} &: UVal_{n+1} \rightarrow (UVal_n \uplus UVal_n) \\ \text{case}_{\rightarrow,n} &: UVal_{n+1} \rightarrow UVal_n \rightarrow UVal_n \\ \text{case}_{\text{unit},n} &\stackrel{\text{def}}{=} \lambda x : UVal_{n+1}. \text{case } x \text{ of } \{ \text{in}_{\text{unit},n} \ x \mapsto x; _ \mapsto \text{omega} \} \\ \text{case}_{\text{bool},n} &\stackrel{\text{def}}{=} \lambda x : UVal_{n+1}. \text{case } x \text{ of } \{ \text{in}_{\text{bool},n} \ x \mapsto x; _ \mapsto \text{omega} \} \\ \text{case}_{\times,n} &\stackrel{\text{def}}{=} \lambda x : UVal_{n+1}. \text{case } x \text{ of } \{ \text{in}_{\times,n} \ x \mapsto x; _ \mapsto \text{omega} \} \\ \text{case}_{\uplus,n} &\stackrel{\text{def}}{=} \lambda x : UVal_{n+1}. \text{case } x \text{ of } \{ \text{in}_{\uplus,n} \ x \mapsto x; _ \mapsto \text{omega} \} \\ \text{case}_{\rightarrow,n} &\stackrel{\text{def}}{=} \lambda x : UVal_{n+1}. \lambda y : UVal_n. \text{case } x \text{ of } \\ & \quad \{ \text{in}_{\rightarrow,n} \ z \mapsto z \ y; _ \mapsto \text{omega} \} \end{aligned}$$

Figure 10. Basic tools for working with $UVal_n$. The subscript of omega is omitted when it is clear from the context.

non-zero levels, the type $UVal_{n+1}$ is a disjunct sum of base values (the second occurrence of Unit and Bool), products and coproducts of $UVal_n$ s and functions mapping a $UVal_n$ to a $UVal_n$. All of these cases are used to emulate a corresponding λ^u value. Additionally, at every level including $n = 0$, the type $UVal_n$ contains a Unit case which is needed to represent an arbitrary λ^u value in cases where the precision of the approximate emulation is insufficient to provide more information. Note that the two occurrences of Unit in the definition of $UVal_{n+1}$ are intentional. The first is used for imprecisely representing arbitrary λ^u terms while the second accurately represents λ^u **unit** values.

To work with $UVal_n$ values, we need basic tools for dealing with sum types: tag injections and case extractions (Fig. 10). Functions $\text{in}_{\text{unk},n}$, $\text{in}_{\text{unit},n}$, $\text{in}_{\text{bool},n}$, $\text{in}_{\times,n}$, $\text{in}_{\uplus,n}$, $\text{in}_{\rightarrow,n}$ are convenient names for nested applications of coproduct injection functions for the nested coproduct in the definition of $UVal_{n+1}$. The term unk_n produces either the single value of $UVal_0$ or uses $\text{in}_{\text{unk},n}$ to produce a $UVal_n$ value representing a 0-precision approximate back-translation of an arbitrary untyped term. For using $UVal_n$ values, we define functions $\text{case}_{\text{unit},n}$, $\text{case}_{\text{bool},n}$, $\text{case}_{\times,n}$, $\text{case}_{\uplus,n}$, $\text{case}_{\rightarrow,n}$ using a somewhat liberal pattern matching syntax that can be easily desugared to nested **case** expressions. The functions are lambdas that inspect their $UVal_{n+1}$ argument and return the contained value if it is in the appropriate branch of the coproduct, or diverge otherwise. To achieve divergence, we use a term omega_τ constructed using **fix**. We simply write omega when the type τ can be inferred from the context.

5.2 λ^u values vs $UVal$

To make the correspondence between a λ^u term and its emulation in $UVal_n$ more exact, this section fills in the definition of $\mathcal{V}[\llbracket \text{EmuLDV}_{n;p} \rrbracket]_\square$, the missing piece of the logical relations of Fig. 7. Recall that the intuition is that $\text{EmuLDV}_{n;p}$ is a token type that is used to relate λ^u terms to their λ^τ emulation of type $UVal_n$. This relation is done up to an approximateness degree, denoted with $p ::= \text{precise} \mid \text{imprecise}$, that is explained below. Intuitively, the previously presented cases of the logical relations define the re-

lation between a λ^τ term and its compilation. The $\mathcal{V}[\llbracket \text{EmulDV}_{n;p} \rrbracket]_{\square}$ case defines the relation between a λ^u term and its UVal_n -typed back-translation, as motivated in Example 5.1.

Example 5.1 (The need of EmulDV). Consider the term $\mathbf{t} \equiv \text{in}_{\text{Bool};1}(\text{true})$. Since UVal_n is a sum type, according to the definition of $\mathcal{V}[\llbracket \tau \uplus \tau' \rrbracket]$, it can be related only to terms that have the same tag. However, for the back-translation we do not want this, we want that term to be related to the \mathbf{t} term that \mathbf{t} approximates (in this case, true).

Type $\text{EmulDV}_{n;p}$ serves the purpose of bridging the syntactic difference, allowing $\text{in}_{\text{Bool};1}(\text{true})$ and true to be related.

Before explaining the definition of the logical relations for $\text{EmulDV}_{n;p}$, we should explain how we elaborate on the approximatness of the correspondence.

Example 5.2. Consider the UVal_6 value $\text{in}_{\times;5}(\text{in}_{\omega;4}(\text{inl unk}_4), \text{unk}_5)$.

This value might be used by the approximate back-translation to represent the λ^u term $(\text{inl}(\text{unit}, \text{true}), \lambda x. x)$. Our $\mathcal{V}[\llbracket \text{EmulDV}_{n;p} \rrbracket]_{\square}$ specification will enforce that terms of the form $\text{in}_{\times;n}(\cdot, \cdot)$ or $\text{in}_{\omega;n}(\text{inl} \cdot)$ represent the corresponding λ^u constructs, but terms unk_4 and unk_5 can represent arbitrary terms (in this case: a pair of base values and a lambda).

The limited size of the type UVal_n sometimes forces us to resort to unk_n values in the back-translation, making it approximate. However, $\mathcal{V}[\llbracket \text{EmulDV}_{n;p} \rrbracket]_{\square}$ does not allow these unk_n values to occur anywhere, because they could compromise the required precision of our approximate back-translation.

In fact, $\mathcal{V}[\llbracket \text{EmulDV}_{n;p} \rrbracket]_{\square}$ provides two different specifications for the occurrence of unk_n , depending on the value of p . The case where $p = \text{imprecise}$ is used when we are proving $\langle\langle \mathbf{e} \rangle\rangle_n \lesssim \mathbf{e}$, which means roughly that termination of $\langle\langle \mathbf{e} \rangle\rangle_n$ in *any* number of steps implies termination of \mathbf{e} . In this case, $\mathcal{V}[\llbracket \text{EmulDV}_{n;p} \rrbracket]_{\square}$ allows unk_n values to occur everywhere in a back-translation term, and they can correspond to arbitrary λ^u terms. These mild requirements on the correspondence of λ^u terms place a large burden on the code in a back-translation $\langle\langle \mathbf{e} \rangle\rangle_n$. This code must be able to deal with unk_n values and produce behaviour for them that approximates the behaviour of \mathbf{e} for the arbitrary values that the unk_n s correspond with. Luckily, when we are proving $\langle\langle \mathbf{e} \rangle\rangle_n \lesssim \mathbf{e}$, we can achieve this by simply making all the functions in our back-translation diverge whenever they try to use a UVal_n value that happens to be an unk_n . This is sufficient because the approximation $\langle\langle \mathbf{e} \rangle\rangle_n \lesssim \mathbf{e}$ trivially holds when $\langle\langle \mathbf{e} \rangle\rangle_n$ diverges: it essentially only requires that \mathbf{e} terminates whenever $\langle\langle \mathbf{e} \rangle\rangle_n$ does, but nothing needs to be shown when the latter diverges.

Example 5.3 (Relatedness with imprecise). Consider the term $\mathbf{t} \equiv \text{in}_{\times;42}(\text{unk}_{42}, \text{unk}_{42})$. This term will be related to $(\mathbf{t}_1, \mathbf{t}_2)$ at type $\text{EmulDV}_{43;\text{imprecise}}$ for any terms \mathbf{t}_1 and \mathbf{t}_2 and in any world.

The increased index 43 is needed because \mathbf{t} is tagged at 42, so we need an additional step to unfold the tagging.

The case when $p = \text{precise}$ specifies where values unk_n are allowed when we are proving that $\langle\langle \mathbf{e} \rangle\rangle_n \gtrsim_n \mathbf{e}$, meaning roughly that termination of \mathbf{e} in less than n steps implies termination of $\langle\langle \mathbf{e} \rangle\rangle_n$. In this case, the requirements on the back-translation correspondence are significantly stronger: unk_n is simply ruled out by the definition of $\mathcal{V}[\llbracket \text{EmulDV}_{n;p} \rrbracket]_{\square}$. That does not mean, however, that unk_n cannot occur inside related terms, rather that unk_n can only occur at depths that cannot be reached using the number of steps in the world.

Example 5.4 (Relatedness with precise). Consider again the term $\mathbf{t} \stackrel{\text{def}}{=} \text{in}_{\times;42}(\text{unk}_{42}, \text{unk}_{42})$. This term will still be related by $\text{EmulDV}_{43;\text{precise}}$ to $\mathbf{t} \stackrel{\text{def}}{=} (\mathbf{t}_1, \mathbf{t}_2)$ for any terms \mathbf{t}_1 and \mathbf{t}_2 , but only in

$$\mathcal{V}[\llbracket \text{EmulDV}_{0;p} \rrbracket]_{\square} \stackrel{\text{def}}{=} \{(\underline{W}, \mathbf{v}, \mathbf{v}) \mid \mathbf{v} = \text{unit} \text{ and } p = \text{imprecise}\}$$

$$\mathcal{V}[\llbracket \text{EmulDV}_{n+1;p} \rrbracket]_{\square} \stackrel{\text{def}}{=} \left\{ (\underline{W}, \mathbf{v}, \mathbf{v}) \mid \begin{array}{l} \mathbf{v} \in \text{oftype}(\text{UVal}_{n+1}) \text{ and one of the following holds:} \\ \left\{ \begin{array}{l} \mathbf{v} = \text{in}_{\text{unk};n} \text{ and } p = \text{imprecise} \\ \exists \mathbf{v}'. \mathbf{v} = \text{in}_{\text{unit};n}(\mathbf{v}') \text{ and } (\underline{W}, \mathbf{v}', \mathbf{v}) \in \mathcal{V}[\llbracket \text{Unit} \rrbracket]_{\square} \\ \exists \mathbf{v}'. \mathbf{v} = \text{in}_{\text{Bool};n}(\mathbf{v}') \text{ and } (\underline{W}, \mathbf{v}', \mathbf{v}) \in \mathcal{V}[\llbracket \text{Bool} \rrbracket]_{\square} \\ \exists \mathbf{v}'. \mathbf{v} = \text{in}_{\times;n}(\mathbf{v}') \text{ and} \\ \quad (\underline{W}, \mathbf{v}', \mathbf{v}) \in \mathcal{V}[\llbracket \text{EmulDV}_{n;p} \times \text{EmulDV}_{n;p} \rrbracket]_{\square} \\ \exists \mathbf{v}'. \mathbf{v} = \text{in}_{\omega;n}(\mathbf{v}') \text{ and} \\ \quad (\underline{W}, \mathbf{v}', \mathbf{v}) \in \mathcal{V}[\llbracket \text{EmulDV}_{n;p} \uplus \text{EmulDV}_{n;p} \rrbracket]_{\square} \\ \exists \mathbf{v}'. \mathbf{v} = \text{in}_{\rightarrow;n}(\mathbf{v}') \text{ and} \\ \quad (\underline{W}, \mathbf{v}', \mathbf{v}) \in \mathcal{V}[\llbracket \text{EmulDV}_{n;p} \rightarrow \text{EmulDV}_{n;p} \rrbracket]_{\square} \end{array} \right\} \end{array} \right\}$$

Figure 11. Specifying the relation between λ^u values and their emulation in $\mathcal{V}[\llbracket \text{EmulDV}_{n;p} \rrbracket]_{\square}$.

$$\begin{array}{ll} \text{toEmul}(\emptyset)_{n;p} = \emptyset & \text{toEmul}(\Gamma, \mathbf{x})_{n;p} = \text{toEmul}(\Gamma)_{n;p}, (\mathbf{x} : \text{EmulDV}_{n;p}) \\ \text{repEmul}(\emptyset) = \emptyset & \text{repEmul}(\Gamma, (\mathbf{x} : \hat{\tau})) = \text{repEmul}(\Gamma), (\mathbf{x} : \text{repEmul}(\hat{\tau})) \\ \text{repEmul}(\tau) = \tau & \text{repEmul}(\text{EmulDV}_{n;p}) = \text{UVal}_n \end{array}$$

Figure 12. Helper functions for $\text{EmulDV}_{n;p}$.

worlds \underline{W} such that $\text{lev}(\underline{W}) = 0$. More precisely, our specification will state that $(\underline{W}, \mathbf{t}, \mathbf{t}) \in \mathcal{V}[\llbracket \text{EmulDV}_{43;\text{precise}} \rrbracket]_{\square}$ iff

$$(\underline{W}, (\text{unk}_{42}, \text{unk}_{42}), \mathbf{t}) \in \mathcal{V}[\llbracket \text{EmulDV}_{42;\text{precise}} \times \text{EmulDV}_{42;\text{precise}} \rrbracket]_{\square}.$$

By the definition in Fig. 7, this requires in turn that $(\underline{W}, \text{unk}_{42}, \mathbf{t}_1)$ and $(\underline{W}, \text{unk}_{42}, \mathbf{t}_2)$ are in $\triangleright \mathcal{V}[\llbracket \text{EmulDV}_{42;\text{precise}} \rrbracket]_{\square}$. However if $\text{lev}(\underline{W}) = 0$, then this is true by definition of the \triangleright operator, independent of the requirements of $\mathcal{V}[\llbracket \text{EmulDV}_{42;\text{precise}} \rrbracket]_{\square}$.

Intuitively, it is sufficient to only forbid unk_n at depths lower than the number of steps left in the world because we are proving $\langle\langle \mathbf{e} \rangle\rangle_n \gtrsim_n \mathbf{e}$ (emphasis on the index n of \gtrsim_n). So, if \mathbf{e} terminates in *less than n steps*, then the evaluation of \mathbf{e} cannot have used values that are deeper than level n in any UVal_n . The corresponding execution of $\langle\langle \mathbf{e} \rangle\rangle_n$ will also not have had a chance to encounter the unk_n s. Therefore, the executions must have behaved identically.

With this approximation aspect explained, Fig. 11 presents the definition of $\mathcal{V}[\llbracket \text{EmulDV}_{n;p} \rrbracket]_{\square}$. For relating terms \mathbf{v} and \mathbf{v} in a world \underline{W} , the definition requires that \mathbf{v} has the right type and that $p = \text{imprecise}$ if \mathbf{v} is unk_n . Additionally, the structure of the λ^τ term stripped of its UVal_n tag and the structure of the λ^u term must coincide. Formally, this is expressed by the following conditions: $(\underline{W}, \mathbf{v}', \mathbf{v})$ are in $\mathcal{V}[\llbracket \mathbf{B} \rrbracket]_{\square}$, $\mathcal{V}[\llbracket \text{EmulDV}_{n;p} \times \text{EmulDV}_{n;p} \rrbracket]_{\square}$, $\mathcal{V}[\llbracket \text{EmulDV}_{n;p} \uplus \text{EmulDV}_{n;p} \rrbracket]_{\square}$ or $\mathcal{V}[\llbracket \text{EmulDV}_{n;p} \rightarrow \text{EmulDV}_{n;p} \rrbracket]_{\square}$ if $\mathbf{v} = \text{in}_{\mathbf{B};n}(\mathbf{v}')$, $\mathbf{v} = \text{in}_{\times;n}(\mathbf{v}')$, $\mathbf{v} = \text{in}_{\omega;n}(\mathbf{v}')$ or $\mathbf{v} = \text{in}_{\rightarrow;n}(\mathbf{v}')$ respectively.

In addition to $\text{EmulDV}_{n;p}$, we still need to define two helper functions (Fig. 12) that were left open in Fig. 7. The first one, $\text{toEmul}(\cdot)_{n;p}$, turns an untyped Γ into one where all variables are mapped to $\text{EmulDV}_{n;p}$. The second one, $\text{repEmul}(\cdot)$, re-maps all variables of a Γ that are of type $\text{EmulDV}_{n;p}$ to type UVal_n .

The adequacy property of the logical relations (Lemma 1) hold for the complete definition of the logical relations, including the definition for $\mathcal{V}[\llbracket \text{EmulDV}_{n;p} \rrbracket]_{\square}$.

$$\begin{aligned}
& \text{downgrade}_{n;d} : \text{UVal}_{n+d} \rightarrow \text{UVal}_n \\
& \text{downgrade}_{0;d} \stackrel{\text{def}}{=} \lambda v : \text{UVal}_d \text{unit} \\
& \text{downgrade}_{n+1;d} \stackrel{\text{def}}{=} \lambda x : \text{UVal}_{n+d+1}. \text{case } x \text{ of} \\
& \left| \begin{array}{l}
\text{in}_{\text{unk},n+d} \mapsto \text{in}_{\text{unk},n} \\
\text{in}_{\text{Unit},n+d} \mathbf{y} \mapsto \text{in}_{\text{Unit},n} \mathbf{y} \\
\text{in}_{\text{Bool},n+d} \mathbf{y} \mapsto \text{in}_{\text{Bool},n} \mathbf{y} \\
\text{in}_{\times,n+d} \mathbf{y} \mapsto \text{in}_{\times,n} (\text{downgrade}_{n;d} \mathbf{y}.1, \text{downgrade}_{n;d} \mathbf{y}.2) \\
\text{in}_{\Psi,n+d} \mathbf{y} \mapsto \text{in}_{\Psi,n} \text{case } \mathbf{y} \text{ of} \left| \begin{array}{l}
\text{inl } \mathbf{x} \mapsto \text{inl} (\text{downgrade}_{n;d} \mathbf{x}) \\
\text{inr } \mathbf{x} \mapsto \text{inr} (\text{downgrade}_{n;d} \mathbf{x})
\end{array} \right. \\
\text{in}_{\rightarrow,n+d} \mathbf{y} \mapsto \text{in}_{\rightarrow,n} (\lambda z : \text{UVal}_n. \text{downgrade}_{n;d} \\
\qquad \qquad \qquad (\mathbf{y} (\text{upgrade}_{n;d} \mathbf{z})))
\end{array} \right. \\
& \text{upgrade}_{n;d} : \text{UVal}_n \rightarrow \text{UVal}_{n+d} \\
& \text{upgrade}_{0;d} \stackrel{\text{def}}{=} \lambda x : \text{UVal}_0. \text{unk}_d \\
& \text{upgrade}_{n+1;d} \stackrel{\text{def}}{=} \lambda x : \text{UVal}_{n+1}. \text{case } x \text{ of} \\
& \left| \begin{array}{l}
\text{in}_{\text{unk},n} \mapsto \text{in}_{\text{unk},n+d}; \\
\text{in}_{\text{Unit},n} \mathbf{y} \mapsto \text{in}_{\text{Unit},n+d} \mathbf{y}; \\
\text{in}_{\text{Bool},n} \mathbf{y} \mapsto \text{in}_{\text{Bool},n+d} \mathbf{y}; \\
\text{in}_{\times,n} \mathbf{y} \mapsto \text{in}_{\times,n+d} (\text{upgrade}_{n;d} \mathbf{y}.1, \text{upgrade}_{n;d} \mathbf{y}.2) \\
\text{in}_{\Psi,n} \mathbf{y} \mapsto \text{in}_{\Psi,n+d} \text{case } \mathbf{y} \text{ of} \left| \begin{array}{l}
\text{inl } \mathbf{x} \mapsto \text{inl} (\text{upgrade}_{n;d} \mathbf{x}) \\
\text{inr } \mathbf{x} \mapsto \text{inr} (\text{upgrade}_{n;d} \mathbf{x})
\end{array} \right. \\
\text{in}_{\rightarrow,n} \mathbf{y} \mapsto \text{in}_{\rightarrow,n+d} (\lambda z : \text{UVal}_n. \text{upgrade}_{n;d} \\
\qquad \qquad \qquad (\mathbf{y} (\text{downgrade}_{n;d} \mathbf{z})))
\end{array} \right.
\end{aligned}$$

Figure 13. Upgrade and downgrade for UVal_n .

5.3 Upgrading and Downgrading Values

Figure 13 defines the functions $\text{downgrade}_{n+d} : \text{UVal}_{n+d} \rightarrow \text{UVal}_n$ and $\text{upgrade}_{n+d} : \text{UVal}_n \rightarrow \text{UVal}_{n+d}$ (by induction on n) that we talked about before. Most cases simply work structurally over the type, but some are more interesting. There is a contravariance in the cases for function values in both downgrade_{n+d} and upgrade_{n+d} : a function $\text{UVal}_n \rightarrow \text{UVal}_n$ is turned into a function of type $\text{UVal}_{n+d} \rightarrow \text{UVal}_{n+d}$ by constructing a wrapper that downgrades the argument and upgrades the result and vice versa. Unknown values are always mapped to unknown values, but additionally, the case for downgrade_{n+d} when $n = 0$ will throw away the information contained in its argument of type UVal_d and simply returns the single unknown value in UVal_0 . Note that downgrade_{n+d} and upgrade_{n+d} are not inverse functions, since $\text{downgrade}_{n+d} \text{upgrade}_{n+d} \mathbf{t}$, the reverse ($\mathbf{t} \equiv \text{upgrade}_{n+d} \text{downgrade}_{n+d} \mathbf{t}$) is not true, since applying downgrade first reduces precision.

Example 5.5 (Downgrading terms). Suppose that we want to emulate a λ^u term $\lambda x. \langle x, x \rangle$ in UVal_n for a sufficiently-large n . We would expect roughly the following λ^τ term:

$$\text{in}_{\rightarrow;n-1} (\lambda x : \text{UVal}_{n-1}. \text{in}_{\times;n-2} \langle \mathbf{x}, \mathbf{x} \rangle)$$

Indices $n-1$ and $n-2$ of the UVal_n constructors are imposed by the well-typedness constraints. However, even this is not enough to guarantee well-typedness. With a closer inspection, the variable \mathbf{x} of type UVal_{n-1} is used where a term of type UVal_{n-2} is required (it is inside a pair tagged with $\text{in}_{\times;n-2}$). This is a problem of type safety, not precision of approximation. Since \mathbf{x} appears inside a pair, inspecting \mathbf{x} for any number of steps requires at least one additional step to first project it out of the pair. In other words,

$$\begin{aligned}
& \text{emulate}_n(\mathbf{t}) : \text{UVal}_n \\
& \text{emulate}_n(\text{unit}) \stackrel{\text{def}}{=} \text{downgrade}_{n+1} (\text{in}_{\text{Unit},n} \text{unit}) \\
& \text{emulate}_n(\text{true}) \stackrel{\text{def}}{=} \text{downgrade}_{n+1} (\text{in}_{\text{Bool},n} \text{true}) \\
& \text{emulate}_n(\text{false}) \stackrel{\text{def}}{=} \text{downgrade}_{n+1} (\text{in}_{\text{Bool},n} \text{false}) \\
& \text{emulate}_n(\mathbf{x}) \stackrel{\text{def}}{=} \mathbf{x} \\
& \text{emulate}_n(\lambda x. \mathbf{t}) \stackrel{\text{def}}{=} \text{downgrade}_{n+1} (\text{in}_{\rightarrow,n} (\lambda x : \text{UVal}_n. \text{emulate}_n(\mathbf{t}))) \\
& \text{emulate}_n(\mathbf{t}_1 \mathbf{t}_2) \stackrel{\text{def}}{=} \text{case}_{\rightarrow,n} (\text{upgrade}_{n+1} (\text{emulate}_n(\mathbf{t}_1))) \text{emulate}_n(\mathbf{t}_2) \\
& \text{emulate}_n(\langle \mathbf{t}_1, \mathbf{t}_2 \rangle) \stackrel{\text{def}}{=} \text{downgrade}_{n+1} (\text{in}_{\times,n} \langle \text{emulate}_n(\mathbf{t}_1), \text{emulate}_n(\mathbf{t}_2) \rangle) \\
& \text{emulate}_n(\text{inl } \mathbf{t}) \stackrel{\text{def}}{=} \text{downgrade}_{n+1} (\text{in}_{\Psi,n} (\text{inl } \text{emulate}_n(\mathbf{t}))) \\
& \text{emulate}_n(\text{inr } \mathbf{t}) \stackrel{\text{def}}{=} \text{downgrade}_{n+1} (\text{in}_{\Psi,n} (\text{inr } \text{emulate}_n(\mathbf{t}))) \\
& \text{emulate}_n(\mathbf{t}.1) \stackrel{\text{def}}{=} (\text{case}_{\times,n} (\text{upgrade}_{n+1} (\text{emulate}_n(\mathbf{t})))) .1 \\
& \text{emulate}_n(\mathbf{t}.2) \stackrel{\text{def}}{=} (\text{case}_{\times,n} (\text{upgrade}_{n+1} (\text{emulate}_n(\mathbf{t})))) .2 \\
& \text{emulate}_n(\mathbf{t}; \mathbf{t}') \stackrel{\text{def}}{=} (\text{case}_{\text{Unit},n} (\text{upgrade}_{n+1} (\text{emulate}_n(\mathbf{t})))) ; \text{emulate}_n(\mathbf{t}') \\
& \text{emulate}_n(\text{wrong}) \stackrel{\text{def}}{=} \text{omega} \\
& \text{emulate}_n(\text{case } \mathbf{t}_1 \text{ of inl } \mathbf{x} \mapsto \mathbf{t}_2 \mid \text{inr } \mathbf{x} \mapsto \mathbf{t}_3) \stackrel{\text{def}}{=} \\
& \quad \text{case } \text{case}_{\Psi,n} (\text{upgrade}_{n+1} (\text{emulate}_n(\mathbf{t}_1))) \text{ of} \\
& \quad \quad \text{inl } \mathbf{x} \mapsto \text{emulate}_n(\mathbf{t}_2) \mid \text{inr } \mathbf{x} \mapsto \text{emulate}_n(\mathbf{t}_3) \\
& \text{emulate}_n(\text{if } \mathbf{t} \text{ then } \mathbf{t}_1 \text{ else } \mathbf{t}_2) \stackrel{\text{def}}{=} \\
& \text{if } (\text{case}_{\text{Bool},n} (\text{upgrade}_{n+1} (\text{emulate}_n(\mathbf{t})))) \text{ then } \text{emulate}_n(\mathbf{t}_1) \text{ else} \\
& \quad \text{emulate}_n(\mathbf{t}_2)
\end{aligned}$$

Figure 14. Emulating λ^u terms in UVal_n .

for the pair to be a precise approximation up to $\leq n-1$ steps, \mathbf{x} needs only to be precise up to $n-2$ steps. It is then safe to throw away one level of precision and downgrade \mathbf{x} from type UVal_{n-1} to UVal_{n-2} .

We will use the function downgrade for the situation of Example 5.5 and similar ones in the next sections. In dual situations we will need to upgrade terms from type UVal_n to UVal_{n+d} . This will neither increase precision of the approximation, nor decrease it.

The correctness property for downgrade and upgrade is stated in the following lemma.

Lemma 6 (Compatibility lemma for upgrade_{n+d} and downgrade_{n+d}). *Suppose that either ($n < m$ and $p = \text{precise}$) or ($\square = \lesssim$ and $p = \text{imprecise}$). Then*

- If $\Gamma \vdash \mathbf{t} \square_n \mathbf{t} : \text{EmulDV}_{m+d;p}$, then $\Gamma \vdash \text{downgrade}_{m+d} \mathbf{t} \square_n \mathbf{t} : \text{EmulDV}_{m;p}$.
- If $\Gamma \vdash \mathbf{t} \square_n \mathbf{t} : \text{EmulDV}_{m;p}$, then $\Gamma \vdash \text{upgrade}_{m+d} \mathbf{t} \square_n \mathbf{t} : \text{EmulDV}_{m+d;p}$.

This lemma covers both situations that we discussed previously. It requires that either $n < m$ (so that the results only hold in worlds \underline{W} with $\text{lev}(\underline{W}) \leq n < m$), in which case $p = \text{precise}$, or $\square = \lesssim$ and $p = \text{imprecise}$. If that is the case, the lemma says that if a term \mathbf{t} is related to \mathbf{t} by $\text{EmulDV}_{m+d;p}$ (or $\text{EmulDV}_{m;p}$) then it stays related to \mathbf{t} after upgrading or downgrading.

5.4 Emulation

Having defined downgrade and upgrade , Fig. 14 defines the emulate_n function. That function maps arbitrary λ^u terms to their approximate back-translation: λ^τ terms of type UVal_n . emulate_n is defined by induction on \mathbf{t} . The different cases follow the same

$$\begin{aligned}
& \text{extract}_{\tau;n} : \text{UVal}_n \rightarrow \tau \\
& \text{extract}_{\tau;0} \stackrel{\text{def}}{=} \lambda x : \text{UVal}_0. \text{omega} \\
& \text{extract}_{\text{Unit};n+1} \stackrel{\text{def}}{=} \lambda x : \text{UVal}_{n+1}. \text{case}_{\text{Unit};n} x \\
& \text{extract}_{\text{Bool};n+1} \stackrel{\text{def}}{=} \lambda x : \text{UVal}_{n+1}. \text{case}_{\text{Bool};n} x \\
& \text{extract}_{\tau_1 \rightarrow \tau_2;n+1} \stackrel{\text{def}}{=} \lambda x : \text{UVal}_{n+1}. \lambda x : \tau_1. \text{extract}_{\tau_2;n} \\
& \quad (\text{case}_{\rightarrow;n} x (\text{inject}_{\tau_1;n} x)) \\
& \text{extract}_{\tau_1 \times \tau_2;n+1} \stackrel{\text{def}}{=} \lambda x : \text{UVal}_{n+1}. (\text{extract}_{\tau_1;n} (\text{case}_{\times;n} x).1, \\
& \quad \text{extract}_{\tau_2;n} (\text{case}_{\times;n} x).2) \\
& \text{extract}_{\tau_1 \uplus \tau_2;n+1} \stackrel{\text{def}}{=} \lambda x : \text{UVal}_{n+1}. \text{case}_{\uplus;n} x \text{ of} \\
& \quad \left| \begin{array}{l} \text{inl } y \rightarrow \text{inl} (\text{extract}_{\tau_1;n} y) \\ \text{inr } y \rightarrow \text{inr} (\text{extract}_{\tau_2;n} y) \end{array} \right. \\
& \text{inject}_{\tau;n} : \tau \rightarrow \text{UVal}_n \\
& \text{inject}_{\tau;0} \stackrel{\text{def}}{=} \lambda x : \tau. \text{omega} \\
& \text{inject}_{\text{Unit};n+1} \stackrel{\text{def}}{=} \lambda x : \text{Unit}. \text{in}_{\text{Unit};n} x \\
& \text{inject}_{\text{Bool};n+1} \stackrel{\text{def}}{=} \lambda x : \text{Bool}. \text{in}_{\text{Bool};n} x \\
& \text{inject}_{\tau_1 \rightarrow \tau_2;n+1} \stackrel{\text{def}}{=} \lambda x : \tau_1 \rightarrow \tau_2. \text{in}_{\rightarrow;n} (\lambda x : \text{UVal}_n. \\
& \quad \text{inject}_{\tau_2;n} (x (\text{extract}_{\tau_1;n} x))) \\
& \text{inject}_{\tau_1 \times \tau_2;n+1} \stackrel{\text{def}}{=} \lambda x : \tau_1 \times \tau_2. \text{in}_{\times;n} (\text{inject}_{\tau_1;n} x.1, \\
& \quad \text{inject}_{\tau_2;n} x.2) \\
& \text{inject}_{\tau_1 \uplus \tau_2;n+1} \stackrel{\text{def}}{=} \lambda x : \tau_1 \uplus \tau_2. \text{in}_{\uplus;n} (\text{case } x \text{ of} \\
& \quad \left| \begin{array}{l} \text{inl } y \mapsto \text{inr} (\text{inject}_{\tau_1;n} y) \\ \text{inr } y \mapsto \text{inr} (\text{inject}_{\tau_2;n} y) \end{array} \right)
\end{aligned}$$

Figure 15. Injecting λ^τ values into UVal_n .

pattern: every term t is mapped to a λ^τ term constructed recursively from the emulation of sub-terms, producing and consuming UVal_n terms wherever t works with untyped terms. Additionally, the definitions use $\text{upgrade}_{n;1}$ and $\text{downgrade}_{n;1}$ to make the resulting term type-check, as discussed in the previous section. For example, the case for pairs applies $\text{in}_{\times;n}$ to a pair constructed from the emulations of its components. Since this produces a UVal_{n+1} , $\text{downgrade}_{n;1}$ is used to downgrade this to a UVal_n term. Finally, the untyped term **wrong** is back-translated to a divergent term.

The back-translation produced by emulate_n is necessarily approximate, as the type UVal_n is not large enough for back-translating arbitrary terms. Inaccuracies in the back-translation are introduced in the calls to $\text{downgrade}_{n;1}$ in several of the cases. The approximation is accurate enough for the following lemma to hold.

Lemma 7 (Emulate relates at EmulDV). *If $\Gamma \vdash t$, and if $(m > n$ and $p = \text{precise})$ or $(\square = \lesssim$ and $p = \text{imprecise})$, then we have that $\text{toEmul}(\Gamma)_{m;p} \vdash \text{emulate}_m(t) \square_n t : \text{EmulDV}_{m;p}$.*

Like Lemma 6, Lemma 7 requires that either $n < m$ (so that the results only hold in worlds \underline{W} with $\text{lev}(\underline{W}) \leq n < m$), in which case $p = \text{precise}$, or $\square = \lesssim$ and $p = \text{imprecise}$. This again covers what we need for the two logical approximations of $\langle\langle \mathcal{C} \rangle\rangle_n$ in Fig. 3. The lemma states that the back-translation of any well-scoped term is related to the term by $\text{EmulDV}_{m;p}$, as intended. An analogous result holds for contexts.

Lemma 8 (Emulated relates contexts at EmulDV). *If $\vdash \mathcal{C} : \Gamma' \rightarrow \Gamma$, if $(m > n$ and $p = \text{precise})$ or $(\square = \lesssim$ and $p = \text{imprecise})$, then $\vdash \text{emulate}_m(\mathcal{C}) \square_n \mathcal{C} : \text{toEmul}(\Gamma')_{m;p}, \text{EmulDV}_{m;p} \rightarrow \text{toEmul}(\Gamma)_{m;p}, \text{EmulDV}_{m;p}$*

5.5 Injection and Extraction of Terms

One final thing is missing to construct a back-translation $\langle\langle \mathcal{C} \rangle\rangle_n$ of an untyped program context \mathcal{C} . While $\text{emulate}_n(\mathcal{C})$ produces a λ^τ context that expects a UVal_n value (just like \mathcal{C} expects an arbitrary λ^u value), the back-translation should accept values of a given type τ (the type of the terms t_1 and t_2 that we are compiling). To bridge this difference, Fig. 15 defines a λ^τ function $\text{inject}_{\tau;n}$ of type $\tau \rightarrow \text{UVal}_n$ which injects values of an arbitrary type τ into UVal_n . We define it mutually recursively with a dual function $\text{extract}_{\tau;n} : \text{UVal}_n \rightarrow \tau$ which is needed for contravariantly converting UVal_n arguments to the appropriate type in the $\text{inject}_{\tau;n}$ case for function types.

Generally, $\text{inject}_{\tau;n}$ converts a value v of type τ to a value of type UVal_n that behaves like the compilation $\llbracket v \rrbracket$. The cases for base values use the related tagging and case (e.g., $\text{in}_{\text{Unit};n}$ and $\text{case}_{\text{Bool};n}$) to achieve this. For pair and sum values, $\text{inject}_{\tau;n}$ and $\text{extract}_{\tau;n}$ simply recurse over the structure of the values, respectively applying $\text{in}_{\times;n}$, $\text{in}_{\uplus;n}$ and $\text{case}_{\times;n}$, $\text{case}_{\uplus;n}$ to construct and destruct UVal_n s of a certain expected form. Note that when UVal_n values do not have the form expected for type τ , then $\text{extract}_{\tau;n}$ will diverge by definition of the $\text{case}_{\dots;n}$ functions. This divergence corresponds to the **wrong** that we get when an untyped context attempts to use λ^u values as pairs, disjunct sum values or base values when those values are of a different form.

For function types, $\text{inject}_{\tau;n}$ and $\text{extract}_{\tau;n}$ produce lambdas that contravariantly extract and inject the argument and covariantly inject and extract the result. Finally, when $n = 0$, then the size of our type is insufficient for $\text{extract}_{\tau;n}$ and $\text{inject}_{\tau;n}$ to accurately perform their intended function. Luckily, to obtain the necessary precision of our approximate back-translation, it is sufficient for them to simply diverge in this case: they simply return **omega** terms of the expected type.

For a value v of type τ , $\text{inject}_{\tau;n}$ will produce a value UVal_n that behaves as the compilation of v , $\llbracket v \rrbracket$. More precisely and more generally, the following lemma states that if a term t is related to a term \hat{t} at type τ (intuitively if t behaves as \hat{t} when used in a way that is valid according to type τ), then $\text{inject}_{\tau;n} t$ behaves as the emulation of $\text{protect}_\tau t$. A dual result about $\text{extract}_{\tau;n}$ and confine_τ states (intuitively) that if a term t behaves as an emulation of value \hat{t} , then $\text{confine}_\tau t$ will behave as $\text{extract}_{\tau;n} \hat{t}$ when used in ways that are valid according to type τ .

Lemma 9 (Inject is protect and extract is confine). *If $(m \geq n$ and $p = \text{precise})$ or $(\square = \lesssim$ and $p = \text{imprecise})$ and if $\hat{\Gamma} \vdash t \square_n \hat{t} : \tau$, then*

$$\hat{\Gamma} \vdash \text{inject}_{\tau;m} t \square_n \text{protect}_\tau t : \text{EmulDV}_{m;p}.$$

If $(m \geq n$ and $p = \text{precise})$ or $(\square = \lesssim$ and $p = \text{imprecise})$ and if $\hat{\Gamma} \vdash t \square_n \hat{t} : \text{EmulDV}_{m;p}$ then

$$\hat{\Gamma} \vdash \text{extract}_{\tau;m} t \square_n \text{confine}_\tau t : \tau.$$

Example 5.6. Consider again Example 4.1. We have that

$$\emptyset \vdash \lambda x : \text{Unit}. x \square \lambda x : x : \text{Unit} \rightarrow \text{Unit}.$$

$\lambda x : \text{Unit}. x$ behaves like $\lambda x. x$, when the latter is used in ways that are valid for a value of type $\text{Unit} \rightarrow \text{Unit}$. Lemma 9 then yields:

$$\emptyset \vdash \text{inject}_{\tau;n} (\lambda x : \text{Unit}. x) \square_n \text{protect}_{\text{Unit} \rightarrow \text{Unit}} (\lambda x. x) : \text{EmulDV}_{m;n}.$$

For n sufficiently large and modulo some simplifications, these terms become:

$$\begin{aligned}
\text{inject}_{\tau;n} (\lambda x : \text{Unit}. x) &= \text{in}_{\rightarrow;n-1} (\lambda x : \text{UVal}_{n-1}. \\
& \quad \text{in}_{\text{Unit};n-2} (\text{case}_{\text{Unit};n-2} x)) \\
\text{protect}_{\text{Unit} \rightarrow \text{Unit}} (\lambda x. x) &= \lambda x. x; \text{unit}
\end{aligned}$$

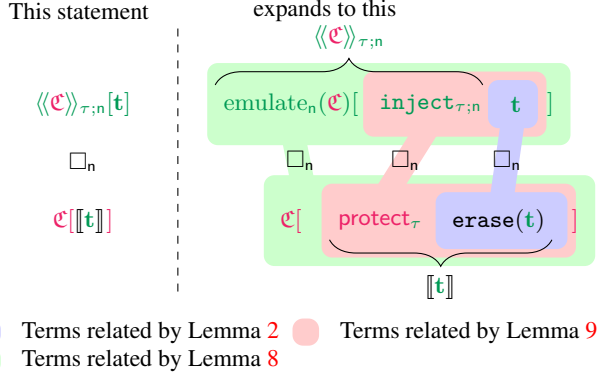


Figure 16. A more accurate picture of related components of compiled term \mathbf{t} , program context \mathcal{C} , compilation $\llbracket \mathbf{t} \rrbracket$ and emulation $\langle\langle \mathcal{C} \rangle\rangle_{\tau;n}$ than in the simplified Fig. 3.

We invite the reader to verify that both expressions behave appropriately when applied to any values \mathbf{v} and \mathbf{v} that are related by $\text{EmulDV}_{n,p}$: for example ($\mathbf{v} = \text{in}_{\text{Unit},n-1} \text{unit}$ and $\mathbf{v} = \text{unit}$), ($\mathbf{v} = \text{in}_{\rightarrow;n-1} (\lambda \mathbf{x} : \text{UVal}_{n-1} \mathbf{x})$ and $\mathbf{v} = \lambda \mathbf{x}. \mathbf{x}$) or ($\mathbf{v} = \text{unk}_n$, \mathbf{v} is any λ^u term and $\square \approx \lesssim$).

5.6 Approximate Back-Translation

We are now ready to define the approximate back-translation $\langle\langle \mathcal{C} \rangle\rangle_{\tau;n}$ of an arbitrary untyped context \mathcal{C} with a hole of type τ . However, before we do, we need to correct a few simplifications that were made in Fig. 3.

First, as we have already explained, the back-translation $\langle\langle \mathcal{C} \rangle\rangle_n$ does not just depend on n but also on the type τ of the terms \mathbf{t}_1 and \mathbf{t}_2 that we are compiling. As such, we define the back-translation with τ as an additional parameter.

Definition 2 (n -approximate back-translation $\langle\langle \cdot \rangle\rangle_{\tau;n}$). *The n -approximate back-translation of a context \mathcal{C} with a hole of type τ is defined as follows. $\langle\langle \mathcal{C} \rangle\rangle_{\tau;n} \stackrel{\text{def}}{=} \text{emulate}_n(\mathcal{C})[\text{inject}_{\tau;n} \cdot]$*

A second simplification in Fig. 3 was the fact that we claimed $\langle\langle \mathcal{C} \rangle\rangle_n \gtrsim_n \mathcal{C}$ and $\langle\langle \mathcal{C} \rangle\rangle_n \lesssim \mathcal{C}$. Fig. 16 shows a more accurate picture of the relations that we have. As we will see in the next section, this more accurate picture still allows us to conclude the facts that $\emptyset \vdash \langle\langle \mathcal{C} \rangle\rangle_{\tau;n}[\mathbf{t}_1] \gtrsim_n \mathcal{C}[\llbracket \mathbf{t}_1 \rrbracket] : \text{EmulDV}_{n,\text{precise}}$ and $\emptyset \vdash \langle\langle \mathcal{C} \rangle\rangle_{\tau;n}[\mathbf{t}_2] \lesssim_{n'} \mathcal{C}[\llbracket \mathbf{t}_2 \rrbracket] : \text{EmulDV}_{n,\text{imprecise}}$ so that the proof goes through unchanged.

The correctness of $\langle\langle \cdot \rangle\rangle_{\tau;n}$ is captured in Lemma 10.

Lemma 10 (Correctness of $\langle\langle \cdot \rangle\rangle_{\tau;n}$). *If $\Gamma \vdash \mathbf{t} \square_n \mathbf{t} : \tau$ then $\Gamma \vdash \langle\langle \mathcal{C} \rangle\rangle_{\tau;n}[\mathbf{t}] \square_n \mathcal{C}[\mathbf{t}] : \tau$.*

6. Compiler Full-Abstraction

This section presents the proof that the compiler $\llbracket \cdot \rrbracket$ is fully-abstract (Theorem 3) in terms of the logical relations of Fig. 7. As previously mentioned, this results in proving compiler correctness (Theorem 1) and compiler security (Theorem 2). As suggested by Fig. 1 in Section 1, the lemmas presented in Section 4 are enough to prove compiler correctness of $\llbracket \cdot \rrbracket$. Dually, as suggested by Fig. 3 in Section 1, the lemmas presented in Section 5 are enough to prove compiler security of $\llbracket \cdot \rrbracket$.

Recall from Definition 1 that $\llbracket \mathbf{t} \rrbracket$ is $\text{protect}_{\tau}(\text{erase}(\mathbf{t}))$.

Theorem 1 ($\llbracket \cdot \rrbracket$ is correct). *If $\emptyset \vdash \mathbf{t}_1 : \tau$ and $\emptyset \vdash \mathbf{t}_2 : \tau$ and $\emptyset \vdash \llbracket \mathbf{t}_1 \rrbracket \simeq_{\text{ctx}} \llbracket \mathbf{t}_2 \rrbracket$, then $\emptyset \vdash \mathbf{t}_1 \simeq_{\text{ctx}} \mathbf{t}_2 : \tau$.*

Proof. Take \mathcal{C} so that $\vdash \mathcal{C} : \emptyset, \tau \rightarrow \emptyset, \tau'$. By definition of \simeq_{ctx} , we need to prove that $\mathcal{C}[\llbracket \mathbf{t}_1 \rrbracket] \Downarrow$ iff $\mathcal{C}[\llbracket \mathbf{t}_2 \rrbracket] \Downarrow$. By symmetry, it suffices to prove the \Rightarrow direction.

So, assume that $\mathcal{C}[\llbracket \mathbf{t}_1 \rrbracket] \Downarrow$. We need to prove that $\mathcal{C}[\llbracket \mathbf{t}_2 \rrbracket] \Downarrow$.

Define $\mathcal{C}' \stackrel{\text{def}}{=} \text{erase}(\mathcal{C})$, Lemma 3 yields $\vdash \mathcal{C}' \square \mathcal{C} : \emptyset, \tau \rightarrow \emptyset, \tau'$. By Lemma 5, we get $\emptyset \vdash \mathbf{t}_1 \square \llbracket \mathbf{t}_1 \rrbracket : \tau$ and $\emptyset \vdash \mathbf{t}_2 \square \llbracket \mathbf{t}_2 \rrbracket : \tau$. By definition of $\vdash \mathcal{C}' \square \mathcal{C} : \emptyset, \tau \rightarrow \emptyset, \tau'$, we get (specifically) that $\emptyset \vdash \mathcal{C}'[\mathbf{t}_1] \gtrsim \mathcal{C}'[\llbracket \mathbf{t}_1 \rrbracket] : \tau'$ and $\emptyset \vdash \mathcal{C}'[\mathbf{t}_2] \lesssim \mathcal{C}'[\llbracket \mathbf{t}_2 \rrbracket] : \tau'$. $\mathcal{C}'[\mathbf{t}_1] \Downarrow$ and $\emptyset \vdash \mathcal{C}'[\mathbf{t}_1] \square \mathcal{C}'[\llbracket \mathbf{t}_1 \rrbracket] : \tau'$ imply $\mathcal{C}'[\llbracket \mathbf{t}_1 \rrbracket] \Downarrow$ by Lemma 1. From $\llbracket \mathbf{t}_1 \rrbracket \simeq_{\text{ctx}} \llbracket \mathbf{t}_2 \rrbracket$ and $\mathcal{C}'[\llbracket \mathbf{t}_1 \rrbracket] \Downarrow$, we get that $\mathcal{C}'[\llbracket \mathbf{t}_2 \rrbracket] \Downarrow$. $\emptyset \vdash \mathcal{C}'[\mathbf{t}_2] \square \mathcal{C}'[\llbracket \mathbf{t}_2 \rrbracket] : \tau'$ and $\mathcal{C}'[\llbracket \mathbf{t}_2 \rrbracket] \Downarrow$ yield $\mathcal{C}[\llbracket \mathbf{t}_2 \rrbracket] \Downarrow$ by Lemma 1. \square

Theorem 2 ($\llbracket \cdot \rrbracket$ is secure). *If $\emptyset \vdash \mathbf{t}_1 : \tau$ and $\emptyset \vdash \mathbf{t}_2 : \tau$ and $\mathbf{t}_1 \simeq_{\text{ctx}} \mathbf{t}_2 : \tau$, then $\llbracket \mathbf{t}_1 \rrbracket \simeq_{\text{ctx}} \llbracket \mathbf{t}_2 \rrbracket$.*

Proof. By definition of \simeq_{ctx} , we need to prove that $\mathcal{C}[\llbracket \mathbf{t}_1 \rrbracket] \Downarrow$ iff $\mathcal{C}[\llbracket \mathbf{t}_2 \rrbracket] \Downarrow$. By symmetry, it suffices to prove the \Rightarrow direction.

Take $\vdash \mathcal{C} : \emptyset \rightarrow \emptyset$ and suppose that $\mathcal{C}[\llbracket \mathbf{t}_1 \rrbracket] \Downarrow$.

We need to show that $\mathcal{C}[\llbracket \mathbf{t}_2 \rrbracket] \Downarrow$.

Take n larger than the number of steps in $\mathcal{C}[\llbracket \mathbf{t}_1 \rrbracket] \Downarrow$.

By Lemma 2, we have that $\emptyset \vdash \mathbf{t}_1 \gtrsim_n \text{erase}(\mathbf{t}_1) : \tau$.

By Lemma 9, we then have (taking $m = n \geq n$) that $\emptyset \vdash \text{inject}_{\tau;n} \mathbf{t}_1 \gtrsim_n \llbracket \mathbf{t}_1 \rrbracket : \text{EmulDV}_{n,\text{precise}}$.

Define $\mathcal{C}' \stackrel{\text{def}}{=} \text{emulate}_n(\mathcal{C})$. From Lemma 8, we get that $\vdash \mathcal{C}' \gtrsim_n \mathcal{C} : \emptyset, \text{EmulDV}_{n,\text{precise}} \rightarrow \emptyset, \text{EmulDV}_{n,\text{precise}}$.

It follows that $\emptyset \vdash \mathcal{C}'[\text{inject}_{\tau;n} \mathbf{t}_1] \gtrsim_n \mathcal{C}'[\llbracket \mathbf{t}_1 \rrbracket] : \text{EmulDV}_{n,\text{precise}}$.

By Lemma 1 applied to $\mathcal{C}'[\llbracket \mathbf{t}_1 \rrbracket] \Downarrow$, and by the choice of n , we have that $\mathcal{C}'[\text{inject}_{\tau;n} \mathbf{t}_1] \Downarrow$.

If we take $\mathcal{C}'' = \mathcal{C}'[\text{inject}_{\tau;n} \cdot]$, then from $\mathbf{t}_1 \simeq_{\text{ctx}} \mathbf{t}_2$ and

$\mathcal{C}'[\text{inject}_{\tau;n} \mathbf{t}_1] \Downarrow$ it follows that $\mathcal{C}'[\text{inject}_{\tau;n} \mathbf{t}_2] \Downarrow$.

Take n' the number of steps in the termination of $\mathcal{C}'[\text{inject}_{\tau;n} \mathbf{t}_2] \Downarrow$.

From Lemma 2 we have that $\emptyset \vdash \mathbf{t}_2 \lesssim_{n'} \text{erase}(\mathbf{t}_2) : \tau$.

By Lemma 9, we have that $\emptyset \vdash \text{inject}_{\tau;n} \mathbf{t}_2 \lesssim_{n'} \llbracket \mathbf{t}_2 \rrbracket : \text{EmulDV}_{n,\text{imprecise}}$.

From Lemma 8, we get that $\vdash \mathcal{C} \lesssim_{n'} \mathcal{C}' : \emptyset, \text{EmulDV}_{n,\text{imprecise}} \rightarrow \emptyset, \text{EmulDV}_{n,\text{imprecise}}$.

It follows that $\emptyset \vdash \mathcal{C}[\text{inject}_{\tau;n} \mathbf{t}_2] \lesssim_{n'} \mathcal{C}[\llbracket \mathbf{t}_2 \rrbracket] : \text{EmulDV}_{n,\text{imprecise}}$.

By Lemma 1, by $\mathcal{C}[\text{inject}_{\tau;n} \mathbf{t}_2] \Downarrow$, and by the choice of n' , we have that $\mathcal{C}[\llbracket \mathbf{t}_2 \rrbracket] \Downarrow$ as required. \square

Theorem 3 ($\llbracket \cdot \rrbracket$ is fully-abstract). *If $\emptyset \vdash \mathbf{t}_1 : \tau$, and $\emptyset \vdash \mathbf{t}_2 : \tau$ then $\mathbf{t}_1 \simeq_{\text{ctx}} \mathbf{t}_2 \iff \llbracket \mathbf{t}_1 \rrbracket \simeq_{\text{ctx}} \llbracket \mathbf{t}_2 \rrbracket$.*

Proof. Theorem 1 provides the \Leftarrow direction while Theorem 2 provides the \Rightarrow one. \square

7. Discussion and Future Work

Our interest in fully-abstract compilation comes from a security perspective. We think that a fully-abstract compiler from realistic source languages to a form of assembly that is efficiently executable by processors has important security applications (combining trusted and untrusted code at the assembly level and compartmentalising applications). Unless targeting typed assembly language [19], a crucial step of a secure compiler is a form of secure type erasure. The contribution of this paper is mostly the proof technique that proves the type erasure step secure. We intend to reuse this proof technique in other settings.

There are a number of important problems that need to be solved in order to develop a realistic fully-abstract compiler. A long-standing open problem is fully-abstract compilation of parametric polymorphism to a form of operational sealing primitives [17, 20, 29]. More concretely, several researchers have developed interesting results about fully-abstract compilation from System F to λ^{seal}

(an untyped lambda calculus with sealing primitives), but a fully-abstract compiler in this setting has so far only been conjectured. We believe that the problem is quite related to the one tackled in this paper. Without providing details (for space reasons), an exact back-translation from λ^{seal} to System F seems possible, but only if we assume a form of generally recursive type constructors of kind $* \rightarrow *$, which we cannot add to System F without causing other problems for the compilation. We conjecture an approximate back-translation is what is needed to provide a fully-abstract compilation in this setting and we hope to confirm it in future work.

In other settings, it is not clear whether it is even possible to construct a fully-abstract compiler. For example, if we add typed, higher-order references to λ^τ and untyped references to λ^u , it is not clear if a fully-abstract compiler can be devised. The problem is essentially to choose a representation for typed references and a way of manipulating them that reconciles a number of requirements: (1) trusted code reading from a reference always produces a type-correct value, (2) trusted code writing a type-correct value to a reference always works, (3) untrusted code should be able to read/write type-correct values from references, (4) dynamic type checks or wrappers may only be added where the context could also choose to fail for other reasons (i.e. not at the time of reading/writing a reference by trusted code), (5) efficiency: we do not want to check the contents of all references every time control is passed from trusted code to the context. Several obvious solutions do not work: representing references as objects with read and write methods violates requirement (4), just checking the contents of a reference when it is received from the context is not enough to guarantee (1) and (2). We intend to explore a solution based on trusted but abstract read/write/alloc methods (using sealing primitives as used for parametric polymorphism) but this remains speculation for the moment.

Another interesting problem when compiling to an assembly language is the enforcement of well-bracketed control flow. The question is essentially how to represent return pointers at the assembly level. Even if we prevent functions from accessing parts of the stack and only give them access to an opaque invokable return pointer, they still have ways to misuse them [23]. Imagine a trusted assembly function f invoking an untrusted g . Additionally, assume that g in turn re-invokes f and f simply re-invokes g again. Now g might attempt to invoke the wrong return pointer, returning on its first invocation without first returning on the second. Such an attack breaks the well-bracketedness of control flow that trusted code may rely on in languages without call/cc primitives [10]. Ahmed and Blume have demonstrated a solution for this problem which exploits parametric polymorphism to enforce the invocation of the correct continuation [4], and it is interesting to see if their work can be reused as an intermediate step on the way to assembly language.

On a technical level, we expect few problems for applying our technique of approximate back-translation to all of these settings. The Hur-Dreyer-inspired cross-language logical relations can be applied in diverse settings including ML and assembly and support references (through Kripke worlds), parametric polymorphism (through quantification over abstract type interpretations as relations) and well-bracketed control flow guarantees (through public/private transitions in the transition systems stored in the worlds). We have also shown in this paper that they can be easily modified to an asymmetric setting.

8. Related Work

Secure compilation through full-abstraction was pioneered by Abadi [1] and successfully applied to many different settings [2, 4, 5, 8, 11, 15, 24, 28, 30]. Recently, Parrow has proven which conditions must hold in the source and target languages in order to provide a fully-abstract compiler between the two [22]. Moreover Gorla and Nestmann have investigated fully-abstract transla-

tions for language expressiveness, concluding that full-abstraction is meaningful when it entails properties such as security [12].

Some secure compilation works prove compiler full-abstraction using logical relations [4, 5, 8, 28, 30]. Ahmed and Blume proved that typed closure conversion and CPS transformation are fully-abstract when compiling from System F and the simply-typed λ -calculus (respectively) to System F [4, 5]. Tse and Zdancewic started a line of work to compile the dependency core calculus of Abadi *et al.* (DCC [3]) into System F, effectively proving that non interference can be encoded with parametricity [30]. They achieve an analogous of fully-abstract compilation where contextual equivalence is replaced with non-interference. Due to an imprecision in their proof, the result of Tse and Zdancewic does not hold; Shikuma and Igarashi refined their result for a weaker form of DCC [28]. A fully-abstract translation from DCC to System F was provided by Bowman and Ahmed [8], and that is the closest work to what presented here. While they prove an important result, we believe the formal machinery adopted by Bowman and Ahmed is heavier than the one presented here. For example, we do not need a new logical relation to prove correctness of the back-translation and can reuse the one for terms. Moreover, all the works that securely compile DCC target System F, so they have a typed target language for their secure compilers while in this paper the target language is untyped.

Many other secure compilation works prove full-abstraction by replacing target-level contextual equivalence with another equivalent equivalence (most times it is trace equivalence or bisimilarity) [2, 11, 15, 24]. These works rely on additional results of the equivalence used for full-abstraction to hold, and this can complicate and lengthen proofs relying on this other technique.

As mentioned in Section 1, the presented proof technique borrows from recent results in compiler correctness [6, 7, 13]. These results build cross-language logical relation based on a common language specification in order to prove compiler correctness. Benton and Hur provided a correct compiler from a call-by-value λ -calculus [6] as well as for System F with recursion to a SECD machine [7]. Hur and Dreyer devised a correct compiler between an idealised ML to assembly [13]. The techniques devised in these works were further developed into Relational Transition Systems (RTS) in order to prove both vertically- and horizontally-composable compiler correctness [14, 21]. A different approach to cross-language relations could have been adopting a Matthews and Findler-style multi-language semantics, where source and target language are combined [18]. For example, Perconti and Ahmed devised a two-step correct compiler for System F with existential and recursive types to typed assembly language using multi-language logical relations [25]. As compiler full-abstraction does scale to multi-pass compilers (i.e., it is vertically composable), there was no necessity to use RTS nor multi-language systems.

9. Conclusion

This paper presented a novel proof technique for proving compiler full-abstraction based on asymmetric, cross-language logical relations. The proof technique revolves around an approximate back-translations from target terms (and contexts) to source terms (and contexts). The back-translation is approximate in the sense that the context generated by the back-translation may diverge when the target-level counterpart would not, but not vice versa. The proof technique is demonstrated for a compiler from a simply-typed λ -calculus without recursive types to the untyped λ -calculus; that compiler is proven to be fully-abstract. Although logical relations have been used for full-abstraction proofs, this is the first usage of cross-language logical relations for compiler full-abstraction targeting an untyped language. We believe the techniques developed in this paper scale to languages with more advanced functionalities

and they can be used to prove compiler full-abstraction in richer settings.

References

- [1] M. Abadi. Protection in programming-language translations. In *Secure Internet programming*, pages 19–34. Springer-Verlag, 1999. ISBN 3-540-66130-1.
- [2] M. Abadi and G. D. Plotkin. On protection by layout randomization. *ACM Transactions on Information and System Security*, 15:8:1–8:29, July 2012. ISSN 1094-9224. doi: 10.1145/2240276.2240279.
- [3] M. Abadi, A. Banerjee, N. Heintze, and J. G. Riecke. A core calculus of dependency. In *Proceedings of the 26th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '99, pages 147–160, New York, NY, USA, 1999. ACM. ISBN 1-58113-095-3. doi: 10.1145/292540.292555.
- [4] A. Ahmed and M. Blume. Typed closure conversion preserves observational equivalence. In *Proceedings of the 13th ACM SIGPLAN International Conference on Functional Programming*, ICFP '08, pages 157–168, New York, NY, USA, 2008. ACM. ISBN 978-1-59593-919-7. doi: 10.1145/1411204.1411227.
- [5] A. Ahmed and M. Blume. An equivalence-preserving CPS translation via multi-language semantics. In *Proceedings of the 16th ACM SIGPLAN International Conference on Functional Programming*, ICFP '11, pages 431–444, New York, NY, USA, 2011. ACM. ISBN 978-1-4503-0865-6. doi: 10.1145/2034773.2034830.
- [6] N. Benton and C.-K. Hur. Biorthogonality, step-indexing and compiler correctness. *SIGPLAN Not.*, 44:97–108, Aug. 2009. ISSN 0362-1340. doi: 10.1145/1631687.1596567.
- [7] N. Benton and C.-k. Hur. Realizability and compositional compiler correctness for a polymorphic language. Technical report, MSR, 2010.
- [8] W. J. Bowman and A. Ahmed. Noninterference for free. In *Proceedings of the 20th ACM SIGPLAN International Conference on Functional Programming*, ICFP '15, New York, NY, USA, 2015. ACM.
- [9] P.-L. Curien. Definability and full abstraction. *Electron. Notes Theor. Comput. Sci.*, 172:301–310, 2007. ISSN 1571-0661. doi: 10.1016/j.entcs.2007.02.011.
- [10] D. Dreyer, G. Neis, and L. Birkedal. The impact of higher-order state and control effects on local relational reasoning. In *International Conference on Functional Programming*, pages 143–156, 2010. doi: 10.1145/1863543.1863566.
- [11] C. Fournet, N. Swamy, J. Chen, P.-E. Dagand, P.-Y. Strub, and B. Livshits. Fully abstract compilation to javascript. In *POPL '13*, pages 371–384, New York, NY, USA, 2013. ACM. ISBN 978-1-4503-1832-7. doi: 10.1145/2429069.2429114.
- [12] D. Gorla and U. Nestman. Full abstraction for expressiveness: History, myths and facts. *Math Struct Comp Science*, 2014.
- [13] C.-K. Hur and D. Dreyer. A kripke logical relation between ml and assembly. In *Principles of Programming Languages*, pages 133–146. ACM, 2011. doi: 10.1145/1926385.1926402.
- [14] C.-K. Hur, D. Dreyer, G. Neis, and V. Vafeiadis. The marriage of bisimulations and kripke logical relations. *SIGPLAN Not.*, 47:59–72, Jan. 2012. ISSN 0362-1340. doi: 10.1145/2103621.2103666.
- [15] R. Jagadeesan, C. Pitcher, J. Rathke, and J. Riely. Local memory via layout randomization. In *Proceedings of the 2011 IEEE 24th Computer Security Foundations Symposium*, CSF '11, pages 161–174, Washington, DC, USA, 2011. IEEE Computer Society. ISBN 978-0-7695-4365-9. doi: 10.1109/CSF.2011.18.
- [16] A. Kennedy. Securing the .NET programming model. *Theor. Comput. Sci.*, 364(3):311–317, Nov. 2006. ISSN 0304-3975. doi: 10.1016/j.tcs.2006.08.014.
- [17] J. Matthews and A. Ahmed. Parametric polymorphism through run-time sealing or, theorems for low, low prices! In *Programming Languages and Systems*, volume 4960 of *Lecture Notes in Computer Science*, pages 16–31. Springer Berlin Heidelberg, 2008. ISBN 978-3-540-78738-9. doi: 10.1007/978-3-540-78739-6_2.
- [18] J. Matthews and R. B. Findler. Operational Semantics for Multi-Language Programs. *ACM Transactions on Programming Languages and Systems*, 31:12:1–12:44, Apr. 2009. ISSN 0164-0925. doi: 10.1145/1498926.1498930.
- [19] G. Morrisett, K. Crary, N. Glew, D. Grossman, R. Samuels, F. Smith, D. Walker, S. Weirich, and S. Zdancewic. TALx86: A Realistic Typed Assembly Language. In *Second Workshop on Compiler Support for System Software*, pages 25–35, 1999.
- [20] G. Neis, D. Dreyer, and A. Rossberg. Non-parametric parametricity. In *Proceedings of the 14th ACM SIGPLAN International Conference on Functional Programming*, pages 135–148. ACM, 2009. ISBN 978-1-60558-332-7. doi: 10.1145/1596550.1596572.
- [21] G. Neis, C.-K. Hur, J.-O. Kaiser, C. McLaughlin, D. Dreyer, and V. Vafeiadis. Pilsner: A Compositionally Verified Compiler for a Higher-Order Imperative Language. In *Proceedings of the 20th ACM SIGPLAN International Conference on Functional Programming*, ICFP '15, New York, NY, USA, 2015. ACM.
- [22] J. Parrow. General conditions for full abstraction. *Math Struct Comp Science*, 2014.
- [23] M. Patrignani. *The Tome of Secure Compilation: Fully Abstract Compilation to Protected Modules Architectures*. PhD thesis, KU Leuven, Leuven, Belgium, May 2015.
- [24] M. Patrignani, P. Agten, R. Strackx, B. Jacobs, D. Clarke, and F. Piessens. Secure Compilation to Protected Module Architectures. *ACM Trans. Program. Lang. Syst.*, 37(2):6:1–6:50, Apr. 2015. ISSN 0164-0925. doi: 10.1145/2699503.
- [25] J. T. Perconti and A. Ahmed. Verifying an open compiler using multi-language semantics. In *ESOP*, volume 8410 of *Lecture Notes in Computer Science*, pages 128–148, 2014.
- [26] B. C. Pierce. *Types and programming languages*. MIT press, 2002.
- [27] G. D. Plotkin. LCF considered as a programming language. *Theoretical Computer Science*, 5:223–255, 1977. doi: http://dx.doi.org/10.1016/0304-3975(77)90044-5.
- [28] N. Shikuma and A. Igarashi. Proving noninterference by a fully complete translation to the simply typed λ -calculus. In M. Okada and I. Satoh, editors, *Advances in Computer Science - ASIAN 2006. Secure Software and Related Issues*, volume 4435 of *Lecture Notes in Computer Science*, pages 301–315. Springer Berlin Heidelberg, 2007. ISBN 978-3-540-77504-1. doi: 10.1007/978-3-540-77505-8_24.
- [29] E. Sumii and B. C. Pierce. A bisimulation for dynamic sealing. *Theor. Comput. Sci.*, 375(1-3):169–192, Apr. 2007.

ISSN 0304-3975. doi: 10.1016/j.tcs.2006.12.032.

- [30] S. Tse and S. Zdancewic. Translating dependency into parametricity. *SIGPLAN Not.*, 39:115–125, Sept. 2004. ISSN 0362-1340. doi: 10.1145/1016848.1016868.