

Adequacy of compositional translations for observational semantics

Manfred Schmidt-Schauß, Joachim Niehren, Jan Schwinghammer, David Sabel

► **To cite this version:**

Manfred Schmidt-Schauß, Joachim Niehren, Jan Schwinghammer, David Sabel. Adequacy of compositional translations for observational semantics. 5th IFIP International Conference on Theoretical Computer Science, Sep 2008, Milano, Italy. inria-00257279v3

HAL Id: inria-00257279

<https://hal.inria.fr/inria-00257279v3>

Submitted on 7 Aug 2008 (v3), last revised 2 Mar 2009 (v7)

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

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

Adequacy of Compositional Translations for Observational Semantics

Manfred Schmidt-Schauß, Joachim Niehren, Jan Schwinghammer, and David Sabel

Abstract We investigate methods and tools for analysing translations between programming languages with respect to observational semantics. The behaviour of programs is observed in terms of may- and must-convergence in arbitrary contexts, and *adequacy* of translations, i.e., the reflection of program equivalence, is taken to be the fundamental correctness condition. For compositional translations we propose a notion of *convergence equivalence* as a means for proving adequacy. This technique avoids explicit reasoning about contexts, and is able to deal with the subtle role of typing in implementations of language extensions.

1 Introduction

Proving correctness of program translations on the basis of operational semantics is an ongoing research topic (see e.g. the recent [?, ?]) that is still poorly understood when it comes to concurrency and mutable state. We are motivated by implementations of language extensions that are often packaged into the language’s library. Typical examples are implementations of channels, buffers, or semaphores using mutable reference cells and futures in Alice ML [?, ?], or using MVars in Concurrent Haskell [?]. Ensuring the correctness of such implementations of higher-level constructs is obviously important.

Manfred Schmidt-Schauß
J. W. Goethe-Universität, Frankfurt, Germany, e-mail: schauss@ki.informatik.uni-frankfurt.de

Joachim Niehren
INRIA, Lille, France, Mostrare Project


Jan Schwinghammer
Saarland University, Saarbrücken, Germany

David Sabel
J. W. Goethe-Universität, Frankfurt, Germany, e-mail: sabel@ki.informatik.uni-frankfurt.de

In this paper we adopt an *observational semantics* based on may- and must-convergence. Two programs are considered equivalent if they exhibit the same may- and must-convergence behaviour in all contexts. This definition is flexible and has been applied to a wide variety of programming languages and calculi in the past. The observation of may- and must-convergence is particularly well-suited for dealing with nondeterminism as it arises in concurrent programming [?, ?, ?].

We study implementations of language extensions in the compilation paradigm, i.e., by viewing them as translations $T : L \rightarrow L'$ from a language L into another language L' . Such translations are usually *compositional* in that $T(C[t]) = T(C)[T(t)]$ for all contexts C and programs t of L . In a naive approach, one might even want to assume that L is a conservative extension of L' so that (non-)equivalences of L' continue to hold in L . However, this fails in many cases (see below) due to subtle typing problems.

A translation $T : L \rightarrow L'$ is *adequate* if $T(s) \sim_{L'} T(t)$ implies $s \sim_L t$ for all programs s and t of L , where \sim_L and $\sim_{L'}$ are the program equivalences of the respective languages. Adequacy is the basic correctness requirement to ensure that program transformations of the target language L' can be soundly applied with respect to observations made in the source language L .



Suppose a translation $T(s)$ is optimized to an equivalent program $s' \sim_{L'} T(s)$ and that s' is the translation of some t , i.e. $T(t) = s'$. Any useful notion of correctness must enforce that s and t are indistinguishable, i.e. $s \sim_L t$. This is precisely what adequacy of T guarantees. With respect to implementations, adequacy opens the possibility of transferring contextual equivalences from the target language L' to the source language L . For non-deterministic and concurrent languages, such equivalences have been established for instance by inductive reasoning using diagram-based methods directly on an underlying small-step operational semantics [?, ?].

Full abstraction extends adequacy by the inverse property, i.e., that program equivalence is also preserved by the translation. In the general situation, however, the language L' may be more expressive than L and allows us to make more distinctions, also on the image $T(L)$. Thus we can have $T(s) \not\sim_{L'} T(t)$ for some expressions s, t with $s \sim_L t$.

In denotational semantics, adequacy and full abstraction are well-studied concepts. In contrast, in this paper we provide a general criterion for proving adequacy of translations that is not tied to specific models. More precisely, we show that convergence equivalence implies adequacy of compositional translations, meaning it is enough to establish that all convergence tests yield the same results before and after the translation. We also provide a criterion for the full abstractness of compositional translations for which the target language is a conservative extension of the source language.

In order to demonstrate these tools, we consider the standard Church encoding of pairs in a call-by-value lambda calculus with a fixed point operator and nondeterministic choice. In order to reason that the encoding of pairs is adequate, one needs to check, for all lambda terms t with pairs and projections, that reduction from t may-converges (must-converges, respectively) if and only if reduction from its encoding $T(t)$ may-converges (must-converges, respectively). However, even in this seemingly well-understood example, this condition *fails* if the lambda calculus is untyped, since the implementation may remove errors, i.e., $T(t)$ terminates more often than t . If the source-language is typed so that stuck expressions are excluded, then our tools apply in a smooth way and show the adequacy of the standard translation, even for differently typed versions of the lambda calculus that is used as target language. Since neither simple typing nor Hindley-Milner polymorphic typing are sufficient to make the source language an extension of the target language, we cannot expect to have an extension situation under type systems that are commonly used in programming languages.

Related work. Various proof methods have been developed for establishing contextual equivalences. These include context lemmas (e.g., [?]), bisimulation methods (for instance, [?]), diagram-based methods (e.g., [?, ?]), and characterizations of contextual equivalence in terms of logical relations (e.g. [?]). In most cases, language extensions and their effect on equivalences are not discussed. There are some notable exceptions: a translation from the core of Standard ML into a typed lambda calculus is given in [?], and full abstraction is shown by exhibiting an inverse mapping, up to contextual equivalence. Adequate translations (with certain additional constraints) between call-by-name and call-by-value versions of PCF are considered in [?], via fully abstract models (necessitating the addition of parallel constructs to the languages) and domain-theoretic techniques. The fact that adequate (and fully abstract) translations compose is exploited in [?], where a syntactic translation is used to lift semantic models for FPC to ones for the lazy lambda calculus. In a similar vein, the recent [?] develops a translation from an aspect-oriented language to an ML-like language, to obtain a model for the former. The adequacy proof follows a similar pattern to ours, but does not abstract away from the particularities of the concrete languages.

Shapiro [?] categorizes implementations and embeddings in concurrent scenarios, but does not provide concrete proof methods based on contextual equivalence. For deterministic languages (where may- and must-convergence agree), frameworks similar to our proposal were considered by Felleisen [?] and Mitchell [?]. Their focus is on comparing languages with respect to their expressive power; the non-deterministic case is only briefly mentioned by Mitchell. Mitchell's work is concerned with (the impossibility of) translations that additionally preserve representation independence of ADTs, and consequently assumes, for the most part, source languages with expressive type systems. Felleisen's work is set in the context of a Scheme-like untyped language. Although the paper discusses the possibility of adding types to get stronger expressiveness statements, the theory of expressiveness is developed by abandoning principles similar to adequacy.

$x, y \in \text{Var}$	$(\beta\text{-CBV}) \quad \mathbb{E}[(\lambda x.t) w] \rightarrow \mathbb{E}[t[w/x]]$
$r, s, t \in \text{Exp}_{cp} ::= w \mid t_1 t_2 \mid t_1 \oplus t_2$	$(\text{FIX}) \quad \mathbb{E}[\mathbf{fix} \lambda x.t] \rightarrow \mathbb{E}[t[(\lambda y.(\mathbf{fix} \lambda x.t)y)/x]]$
$v, w \in \text{Val}_{cp} ::= x \mid \lambda x.t \mid \mathbf{unit} \mid \mathbf{fix}$ $\quad \quad \quad \mid (w_1, w_2) \mid \mathbf{fst} \mid \mathbf{snd}$	$(\oplus\text{L}) \quad \mathbb{E}[w_1 \oplus w_2] \rightarrow \mathbb{E}[w_1]$
	$(\oplus\text{R}) \quad \mathbb{E}[w_1 \oplus w_2] \rightarrow \mathbb{E}[w_2]$

Fig. 1 Syntax of λ_{cp}

$\mathbb{E} ::= [] \mid \mathbb{E}t \mid w\mathbb{E} \mid \mathbb{E} \oplus t \mid w \oplus \mathbb{E}$	$(\text{SEL-F}) \quad \mathbb{E}[\mathbf{fst}(w_1, w_2)] \rightarrow \mathbb{E}[w_1]$
	$(\text{SEL-S}) \quad \mathbb{E}[\mathbf{snd}(w_1, w_2)] \rightarrow \mathbb{E}[w_2]$

Fig. 2 Evaluation Contexts \mathbb{E} **Fig. 3** Small-Step Reduction

Outline. Section 2 recalls the encoding of pairs in the non-deterministic lambda calculus, introduces rigorous notions of observables, and illustrates the need for types. In Section 3 a general framework for proving observational correctness as well as adequacy of translations is introduced. Section 4 shows the adequacy of the pair encoding using a simple type system and discusses two extensions.

2 Non-deterministic Call-by-Value Lambda Calculi

In this section, we recall the call-by-value lambda calculus with a fixed point operator and nondeterministic choice, and present its observational semantics on the basis of may- and must-convergence. We illustrate why Church's encoding of pairs in this calculus fails to be observationally correct in the untyped case.

2.1 Languages

The calculus λ_{cp} is the usual call-by-value lambda calculus extended by a (demonic, see [?]) choice operator, a call-by-value fixed point operator for recursion, pairs (w_1, w_2) and selectors **fst** and **snd** as data structure, and a constant **unit**. Fixing a set of variables Var , the syntax of expressions Exp_{cp} and values Val_{cp} is shown in Fig. 1. The subcalculus λ_c is the calculus without pairs and selectors and will be used as target language. We use Exp_c (Val_c , resp.) for the set of λ_c -expressions (λ_c -values, resp.).

A *context* C is an expression with a hole denoted with $[]$, $C[s]$ is the result of placing the expression s in the hole of C . For both calculi we require call-by-value evaluation contexts \mathbb{E} which are introduced in Fig. 2. With $s_1[s_2/x]$ we denote the capture-free substitution of variable x with s_2 for all free occurrences of x in s_1 . To ease reasoning we assume that the distinct variable convention holds for all expressions, i.e. that the bound variables of an expression are all distinct and free variables are distinct from bound variables.

The reduction rules for both calculi are defined in Fig. 3. Small step reduction \rightarrow_{cp} of λ_{cp} is the union of all six rules, and small step reduction \rightarrow_c of λ_c is the union of the first four rules. We assume that reduction preserves the distinct variable convention by implicitly performing α -renaming if necessary.

2.2 Contextual Equivalence

Let Exp be a language, let $Val \subseteq Exp$ be a set of values and \rightarrow be a reduction relation. Then *may-convergence* for expressions $s \in Exp$ is defined as $s \downarrow$ iff $\exists v \in Val : s \xrightarrow{*} v$, and *must-convergence* is defined as $s \Downarrow$ iff $\forall s' : s \xrightarrow{*} s' \implies s' \downarrow$. For a discussion and motivations for the latter notion see [?, ?, ?]. Note that there is also another notion of must-convergence found in the literature (e.g. [?]), which holds if an expression has only evaluations to values, in particular, if the expression has no infinite evaluations (i.e. if $s \not\rightarrow^\omega$).

For an expression s we also write $s \uparrow$ if $s \downarrow$ does not hold, and say that s is *must-divergent*. We write $s \uparrow$ if s is not must-convergent and then say s is *may-divergent*. Note that may-divergence can equivalently be defined as $s \uparrow$ iff $\exists s' \in Exp : s \xrightarrow{*} s'$ and $s' \uparrow$. This view allows us to use inductive proofs for showing may-divergences. For Exp_c, Val_c , and \rightarrow_c we use \downarrow_c for may-convergence and \Downarrow_c for must-convergence. Accordingly for Exp_{cp}, Val_{cp} , and \rightarrow_{cp} we use \downarrow_{cp} and \Downarrow_{cp} for the predicates.

Contextual equivalence for a (non-deterministic) calculus (Exp, Val, \rightarrow) is defined by observing may- and must-convergence in all contexts. We first define two preorders for both predicates:

$$s_1 \leq_{\downarrow} s_2 \text{ iff } \forall C : C[s_1] \downarrow \implies C[s_2] \downarrow \qquad s_1 \leq_{\Downarrow} s_2 \text{ iff } \forall C : C[s_1] \Downarrow \implies C[s_2] \Downarrow$$

These are combined to obtain the contextual preorder \leq as their intersection $\leq_{\downarrow} \cap \leq_{\Downarrow}$, and the contextual equivalence \sim as $\leq \cap \geq$. To distinguish between the relations for λ_c and λ_{cp} , we index the symbols for the preorders and equivalence with c or cp , respectively, e.g. contextual equivalence in λ_c is \sim_c , and contextual preorder in λ_{cp} is \leq_{cp} .

2.3 Implementation of Pairs

We will mainly investigate the translation enc of λ_{cp} into λ_c as defined in Fig. 4 under different restrictions. Conversely, it is trivial to encode λ_c into λ_{cp} via the identity $inc(s) = s$ (which is more an *embedding* than a translation).

The following counter example shows that the implementation of pairs is not correct in the untyped setting.

$$\begin{array}{ll}
enc(x) & = x & enc(\mathbf{fix}) & = \mathbf{fix} \\
enc(\mathbf{unit}) & = \mathbf{unit} & enc((w_1, w_2)) & = \lambda s. (s \ enc(w_1) \ enc(w_2)) \\
enc(\lambda x.t) & = \lambda x. enc(t) & enc(\mathbf{fst}) & = \lambda p. (p \ \lambda x. \lambda y. x) \\
enc(t_1 \ t_2) & = enc(t_1) \ enc(t_2) & enc(\mathbf{snd}) & = \lambda p. (p \ \lambda x. \lambda y. y) \\
enc(t_1 \oplus t_2) & = enc(t_1) \oplus enc(t_2) & &
\end{array}$$

Fig. 4 Translation of λ_{cp} into λ_c

Example 2.1. Let $t := \mathbf{fst}(\lambda z.z)$. Then $t \uparrow_{cp}$, since t is irreducible and not a value. However, the translation $enc(t)$ results in the expression $t' := (\lambda p.p \ (\lambda x.\lambda y.x)) \ (\lambda z.z)$, which deterministically reduces by some (β -CBV)-reductions to $\lambda x.\lambda y.x$, hence $enc(t) \downarrow_c$. This is clearly not a correct translation, since it removes an error. Therefore, the observations are not preserved by this translation. This example also invalidates the implication $T(p_1) \leq_c T(p_2) \implies p_1 \leq_{cp} p_2$, since $enc(t') = t'$, and hence $enc(t') = t' \leq_c t' = enc(t)$, but $t' \not\leq_{cp} t$ by the arguments above. In the terminology of Definition 3.2 below, the translation enc is not adequate.

This counter example is also valid for deterministic calculi, where may- and must-convergence coincide. There, it is possible to circumvent the problem by weakening the definition of correctness to only one direction of the logical equivalence, $s \downarrow \implies T(s) \downarrow$, but this results in weaker properties and is not the appropriate notion for compilations. In particular, this notion of correctness of a translation (which is called *weak expressibility* in [?]) implies the correctness of a trivial translation that maps all expressions to a (may-) convergent expression.

One potential remedy to the failure of the untyped approach to correctness of translations is to distinguish divergence from typing errors. From a different point of view, this simply means that only correctly typed programs should be considered by a translation: in Section 4.1 we will obtain adequacy after adding a type system to λ_{cp} .

3 Adequacy of Translations

We present a general framework for reasoning about different notions of language translations which are related to correctness.

We assume that languages come equipped with a small-step operational semantics and a notion of observables, expressed through convergence tests, with respect to which contextual equivalence can be defined. Since we are interested in concurrent calculi, a typical case will be the observations of may- and must-termination behavior, as introduced in the previous section. In the following we generalize slightly and, instead of contexts, speak of observers: this makes it easier to fit formalisms without an obvious notion of context into the framework, like abstract machines.

Definition 3.1. A program calculus with observational semantics (*OSP-calculus*) consists of the following components:

- A set \mathcal{T} of *types*, ranged over by τ .
- For every type τ , a set \mathcal{P}_τ of *programs*, ranged over by p .
- For every pair τ_1, τ_2 of types, a set of functions $\mathcal{O}_{\tau_1, \tau_2}$ with $O : \mathcal{P}_{\tau_1} \rightarrow \mathcal{P}_{\tau_2}$ for $O \in \mathcal{O}_{\tau_1, \tau_2}$, called *observers*, such that also the identity function Id_τ is included in $\mathcal{O}_{\tau, \tau}$ for every type τ , and such that $\bigcup_{\tau_1, \tau_2 \in \mathcal{T}} \mathcal{O}_{\tau_1, \tau_2}$ is closed under function composition whenever the types are appropriate.
- A set $\{\Downarrow_1, \dots, \Downarrow_n\}$ of *convergence tests* with $\Downarrow_i : \bigcup_{\tau \in \mathcal{T}} \mathcal{P}_\tau \rightarrow \{\text{true}, \text{false}\}$ for all $i = 1, \dots, n$.

This definition is also applicable to the special case of deterministic calculi, where usually only a single termination predicate is considered. Moreover, it allows for untyped calculi like λ_{cp} by considering a single, ‘universal’ type. The calculus λ_{cp} then fits this definition of OSP-calculus, after identifying a context C with the map $t \mapsto C[t]$, and taking $\{\Downarrow_1, \Downarrow_2\} = \{\downarrow_{cp}, \downarrow_{cp}\}$.

Since this framework has arbitrary observers (not only contexts) and there are types, the observational preorders at type τ are defined as follows, where $p_1, p_2 \in \mathcal{P}_\tau$:

- $p_1 \leq_{\Downarrow_i, \tau} p_2$ iff for all $\tau' \in \mathcal{T}$ and all $O : \tau \rightarrow \tau'$, $O(p_1) \Downarrow_i$ implies $O(p_2) \Downarrow_i$.
- $p_1 \leq_\tau p_2$ iff $\forall i : p_1 \leq_{\Downarrow_i, \tau} p_2$.
- $p_1 \sim_\tau p_2$ iff $p_1 \leq_\tau p_2$ and $p_2 \leq_\tau p_1$.

The relations $\leq_{\Downarrow_i, \tau}$ and \leq_τ are *precongruences*, i.e. they are preorders, and $p_1 \leq_{\Downarrow_i, \tau} p_2$ implies $\forall O : \tau \rightarrow \tau' : O(p_1) \leq_{\Downarrow_i, \tau'} O(p_2)$. For proving the latter implication let O' be an observer with $O'(O(p_1)) \Downarrow_i$. Then $O' \circ O$ is also an observer, hence $O' \circ O(p_2) \Downarrow_i$. Obviously, the same holds for \leq_τ . The relation \sim_τ is a *congruence*, i.e. it is a precongruence and an equivalence relation.

In the following we only consider translations between OSP-calculi that have the same number n of convergence tests $\{\Downarrow_1, \dots, \Downarrow_n\}$, in a fixed ordering. We define some characterizing notions of translations. In the remainder of this section we exhibit their dependencies and prove some consequences.

Definition 3.2. A translation $T : \mathcal{C} \rightarrow \mathcal{C}'$ between two calculi $\mathcal{C} = (\mathcal{T}, \mathcal{P}, \mathcal{O}, \leq)$ and $\mathcal{C}' = (\mathcal{T}', \mathcal{P}', \mathcal{O}', \leq')$ maps types to types $T : \mathcal{T} \rightarrow \mathcal{T}'$, programs to programs $T : \mathcal{P}_\tau \rightarrow \mathcal{P}'_{T(\tau)}$, and observers to observers $T : \mathcal{O}_{\tau, \tau'} \rightarrow \mathcal{O}'_{T(\tau), T(\tau')}$ such that their types correspond for all $\tau, \tau' \in \mathcal{T}$ and such that $T(Id_\tau) = Id_{T(\tau)}$ for all τ .

Adequacy. A translation T is *adequate* iff for all τ , and $p_1, p_2 \in \mathcal{P}_\tau$, $T(p_1) \leq'_{T(\tau)} T(p_2) \implies p_1 \leq_\tau p_2$.

Full abstraction. A translation T is *fully abstract* iff for all τ , and $p_1, p_2 \in \mathcal{P}_\tau$, $p_1 \leq_\tau p_2 \iff T(p_1) \leq'_{T(\tau)} T(p_2)$.

Observational correctness. A translation T is *observationally correct* iff for all $\tau, p \in \mathcal{P}_\tau, O \in \mathcal{O}_{\tau, \tau'}$ and all i : $O(p) \Downarrow_i$ if and only if $T(O)(T(p)) \Downarrow'_i$.

Convergence equivalence. A translation T is *convergence equivalent* (i.e. *preserves and reflects convergence*) iff for all p and convergence tests \Downarrow_i : $p \Downarrow_i$ if and only if $T(p) \Downarrow'_i$.

Compositionality. A translation T is *compositional* iff for all types $\tau, \tau' \in \mathcal{T}$, for all observers $O \in \mathcal{O}_{\tau, \tau'}$ and all programs $p \in \mathcal{P}_{\tau}$ we have $T(O(p)) = T(O)(T(p))$.

If in the following types are omitted, we implicitly assume that type information follows from the context.

As motivated in the Introduction, we consider adequacy as the right notion of correctness. Observational correctness is a sufficient criterion for adequacy (see Proposition 3.3). Convergence equivalence is implied by observational correctness, since T preserves identity observers. For compositional translations, the converse is true, i.e., it is sufficient to prove convergence equivalence in order to prove observational correctness. Full abstraction is not necessary for the adequacy of translations. If it holds in addition, for surjective translations it means that both program calculi are identical w.r.t. \leq .

Note that Definition 3.2 is stated only in terms of convergence tests and sets of observers, and hence only relying on the syntax and the operational semantics. Thus it can be used in all calculi with such a description. In the case of two calculi with convergence tests defined in terms of a small-step semantics, the definition also allows for reduction sequences in the translation that may lead outside of the image of the translation, i.e., that may not be retranslatable.

Proposition 3.3. *For a translation T the following hold:*

1. *If T is compositional, then T is convergence equivalent if and only if T is observationally correct.*
2. *If T is observationally correct, then T is adequate.*

Proof. 1. The only if direction holds, since T preserves identity observers: $Id_{\tau}(p) \Downarrow_i \iff T(Id_{\tau})T(p) \Downarrow'_i \iff Id_{T(\tau)}T(p) \Downarrow'_i \iff T(p) \Downarrow'_i$.

For the if-direction let us assume that T is compositional and convergence equivalent. If $O(p) \Downarrow_i$, then preservation of convergence yields $T(O(p)) \Downarrow'_i$. Compositionality implies $T(O(p)) = T(O)(T(p))$, hence $T(O)(T(p)) \Downarrow'_i$. If $T(O)(T(p)) \Downarrow'_i$ then compositionality implies $T(O(p)) \Downarrow'_i$ so that reflection of convergence yields $O(p) \Downarrow_i$.

2. To show adequacy, let us assume that $T(p_1) \leq_{T(\tau)} T(p_2)$. We must prove that $p_1 \leq_{\tau} p_2$. Thus let O be such that $O(p_1) \Downarrow_i$. By observational correctness this implies $T(O)(T(p_1)) \Downarrow'_i$. From $T(p_1) \leq_{T(\tau)} T(p_2)$, we obtain $T(O)(T(p_2)) \Downarrow'_i$, since $T(O)$ is an admissible observer. Observational correctness in the other direction implies $O(p_2) \Downarrow_i$. This proves $p_1 \leq_{\tau} p_2$. \square

As the following counter examples show, convergence equivalence is in general not sufficient for adequacy, and full abstraction is not implied by observational correctness. Similarly, convergence equivalence is not even implied by full abstraction (and thus neither by adequacy):

Example 3.4 (Convergence equivalence does not imply adequacy). Let the OSP-calculus L have three programs: a, b, c with $a \uparrow$, $b \downarrow$ and $c \downarrow$. Assume there are

two observers O_1, O_2 with $O_1(x) = x$ and $O_2(a) = a, O_2(b) = a, O_2(c) = c$. Then $b \not\sim_L c$. The language L' has three programs A, B, C with $A \uparrow, B \downarrow$ and $C \downarrow$. There is only the identity observer O in L' . Then $B \sim_{L'} C$. Let the translation be defined as $T : L \rightarrow L'$ with $T(a) = A, T(b) = B, T(c) = C$, and $T(O_1) = T(O_2) = O$. Then convergence equivalence holds, but neither equational adequacy nor observational correctness. Note that T is not compositional, since $T(O_2(b)) = A$ while $T(O_2)(T(b)) = O(B) = B$.

Example 3.5 (Observational correctness does not imply full abstraction). A simple example taken from [?] is the identity encoding from the OSP-calculus λ_{cp} without the projections **fst** and **snd** into full λ_{cp} . Then, in the restricted OSP-calculus, all pairs are indistinguishable but the presence of the observers (here simply taken as contexts) **fst** $[\cdot]$ and **snd** $[\cdot]$ in λ_{cp} permits more distinctions to be made.

Example 3.6 (Convergence equivalence is not implied by full abstraction). A trivial example is given by two calculi \mathcal{C} with $p \Downarrow$ for all p , and \mathcal{C}' with the same programs and $\neg p \Downarrow'$ for all p . For the translation $T(p) = p$ for all p it is clear that $\forall p_1, p_2 : p_1 \leq p_2 \iff T(p_1) \leq' T(p_2)$ holds, but T does not preserve convergence.


By standard arguments it can be shown that translations compose:

Proposition 3.7. *Let $\mathcal{C}, \mathcal{C}', \mathcal{C}''$ be program calculi, and $T : \mathcal{C} \rightarrow \mathcal{C}', T' : \mathcal{C}' \rightarrow \mathcal{C}''$ be translations. Then $T' \circ T : \mathcal{C} \rightarrow \mathcal{C}''$ is also a translation, and for every property P from Definition 3.2, if T, T' have property P , then also the composition $T' \circ T$.*

We now consider the case that only new language primitives are added to a language, together with their operational semantics, which are then encoded by the translation. This is usually known as removing ‘syntactic sugar’.

Definition 3.8. An OSP-calculus \mathcal{C} is an *extension* of the OSP-calculus \mathcal{C}' iff there is a compositional translation $\iota : \mathcal{C}' \rightarrow \mathcal{C}$, called an *embedding*, which is injective on the expressions, types and observers, and is convergence equivalent.

Informally, this can be described (after identifying \mathcal{C}' -programs with their image under ι) as follows: every \mathcal{C}' -type is also a \mathcal{C} -type, $\mathcal{P}'_\tau \subseteq \mathcal{P}_\tau$, and $\mathcal{O}'_{\tau, \tau'}$ is a subset of $\mathcal{O}_{\tau, \tau'}$, and the test-predicates coincide on \mathcal{C}' -programs. The embedding of $\mathcal{O}'_{\tau, \tau'}$ into $\mathcal{O}_{\tau, \tau'}$ is slightly more involved, since the \mathcal{C}' -observers are restrictions (as functions) of \mathcal{C} -observers. Note that for the case of contexts as observers, the embedding of $\mathcal{O}'_{\tau, \tau'}$ into $\mathcal{O}_{\tau, \tau'}$ is unique. The conditions imply that an embedding ι is adequate, but not necessarily fully abstract.



If \mathcal{C} is an extension of \mathcal{C}' , then an observationally correct translation $T : \mathcal{C} \rightarrow \mathcal{C}'$ (plus some obvious conditions) has the nice consequence of T and ι being fully abstract.

An example for an embedding is the trivial embedding $inc : \lambda_c \rightarrow \lambda_{cp}$, which is adequate by Proposition 3.3, since the embedding inc is compositional and convergence equivalent. This allows us to reason about contextual equivalence in λ_{cp} and transfer this result to λ_c , i.e. a proof of $t_1 \sim_{cp} t_2$ where t_1, t_2 are also expressions of λ_c directly shows $t_1 \sim_c t_2$. Disproving an equivalence in λ_{cp} , however, does *not* imply that this equivalence is false in λ_c .

Proposition 3.9 (Full Abstraction for Extensions). *Let \mathcal{C} be an extension of \mathcal{C}' , and let $T : \mathcal{C} \rightarrow \mathcal{C}'$ be an observationally correct translation, such that $T \circ \iota$ is the identity on \mathcal{C}' -programs and on \mathcal{C}' -types. Then the translation T as well as the embedding ι are fully abstract.*

Proof. First we show that $T \circ \iota$ is also the identity on \mathcal{C}' -observers: For arbitrary p , the equation $(T \circ \iota)(O)(p) = (T \circ \iota)(O)((T \circ \iota)(p)) = (T \circ \iota)(O(p)) = O(p)$ holds due to the assumption that $T \circ \iota$ is the identity on programs and due to compositionality of $T \circ \iota$. Hence $T \circ \iota$ is also the identity on observers.

We show that T is fully abstract. Adequacy follows from Proposition 3.3. It remains to show the converse condition for full abstraction. Let p_1, p_2 be \mathcal{C} -programs of type τ , and assume $p_1 \leq_{\Downarrow_i, \tau} p_2$. We have to show that $T(p_1) \leq'_{\Downarrow_i, T(\tau)} T(p_2)$. Let O' be a \mathcal{C}' -observer with $O'(T(p_1)) \Downarrow'_i$. Then by definition of ι there exists an observer O of \mathcal{C} with $O := \iota(O')$. Since $T \circ \iota$ is the identity, we have $T(O) = O'$ and thus we obtain $T(O)(T(p_1)) \Downarrow'_i$. Observational correctness implies that $O(p_1) \Downarrow_i$. From $p_1 \leq_{\Downarrow_i, \tau} p_2$ we now derive $O(p_2) \Downarrow_i$. Again observational correctness can be applied and shows that $T(O)(T(p_2)) \Downarrow'_i$. This is equivalent to $O'(T(p_2)) \Downarrow'_i$. Since the observer O' was chosen arbitrarily, we have $T(p_1) \leq'_{\Downarrow_i, T(\tau)} T(p_2)$.

The embedding ι is already shown to be adequate. The missing direction, i.e. that $\iota(p_1) \leq'_{\Downarrow_i, T(\tau)} \iota(p_2)$ implies $p_1 \leq'_{\Downarrow_i, \tau} p_2$ follows from full abstraction of T and the assumption that $T \circ \iota$ is the identity. \square

We also require an generalisation of the extension theorem, since some translations do not satisfy the strong precondition that $T \circ \iota$ is the identity. Therefore, the equivalence \sim_τ on expressions is extended to observers as follows: two observers $O_1, O_2 \in \mathcal{O}_{\tau, \tau'}$ are equivalent, i.e. $O_1 \sim_{\tau'} O_2$, iff for all objects p of type τ , we have $O_1(p) \sim_{\tau'} O_2(p)$. Note that since the \sim_τ are congruences, we also have that $s \sim_\tau t$ and $O_1 \sim_{\tau'} O_2$ imply $O_1(s) \sim_{\tau'} O_2(t)$.

Proposition 3.10 (Full Abstraction for Extensions, Generalized). *Let \mathcal{C} be an extension of \mathcal{C}' , and let $T : \mathcal{C} \rightarrow \mathcal{C}'$ be an observationally correct translation, such that $(T \circ \iota)(p) \sim p$ for all \mathcal{C}' -programs p , and $T \circ \iota$ is the identity on \mathcal{C}' -types. Then the translation T as well as the embedding ι are fully abstract.*

Proof. First we show that $(T \circ \iota)(O) \sim O$ for all \mathcal{C}' -observers O : For arbitrary p , the equation $(T \circ \iota)(O)(p) \sim (T \circ \iota)(O)((T \circ \iota)(p)) = (T \circ \iota)(O(p)) \sim O(p)$ holds

due to the assumption that $T \circ \iota$ is the identity on programs modulo \sim and due to compositionality of $T \circ \iota$. Hence $(T \circ \iota)$ is also the identity modulo \sim on observers.

We show that T is fully abstract. Adequacy follows from Proposition 3.3. It remains to show the converse condition for full abstraction. Let p_1, p_2 be \mathcal{C} -programs of type τ , and assume $p_1 \leq_{\Downarrow_i, \tau} p_2$. We have to show that $T(p_1) \leq'_{\Downarrow_i, T(\tau)} T(p_2)$. Let O' be a \mathcal{C}' -observer with $O'(T(p_1)) \Downarrow'_i$. Then by definition of ι there exists an observer O of \mathcal{C} with $O := \iota(O')$. By the precondition on $T \circ \iota$, we have $T(O) \sim O'$ and thus we obtain $T(O)(T(p_1)) \Downarrow'_i$. Observational correctness implies that $O(p_1) \Downarrow_i$. From $p_1 \leq_{\Downarrow_i, \tau} p_2$ we now derive $O(p_2) \Downarrow_i$. Again observational correctness can be applied and shows that $T(O)(T(p_2)) \Downarrow'_i$. This is equivalent to $O'(T(p_2)) \Downarrow'_i$. Since the observer O' was chosen arbitrarily, we have $T(p_1) \leq'_{\Downarrow_i, T(\tau)} T(p_2)$.

The embedding ι is already shown to be adequate. The missing direction, i.e. that $\iota(p_1) \leq'_{\Downarrow_i, T(\tau)} \iota(p_2)$ implies $p_1 \leq'_{\Downarrow_i, \tau} p_2$ follows from full abstraction of T and the preconditions on $T \circ \iota$ which enforces that $p_1 \sim T(\iota(p_1)) \leq'_{\Downarrow_i, T(\tau)} T(\iota(p_2)) \sim p_2$. \square

Remark 3.11. We can generalize the notion of compositionality of a translation T to *compositionality modulo \sim* , meaning that $T(O(p)) \sim T(O)T(p)$ for all O and p . Appropriate versions of the results of the previous Lemmas, Theorems and Propositions hold: Proposition 3.7 is also valid for compositionality modulo \sim , and Proposition 3.3, Proposition 3.9, and Proposition 3.10 are valid also if compositionality is replaced by compositionality modulo \sim .

4 Adequacy of Pair Encoding

We analyze the translation enc on the untyped language λ_c . Inspecting the definition of enc the following lemma is easy to verify:

Lemma 4.1. *For all $s \in \lambda_{cp}$: s is a λ_{cp} -value iff $enc(s)$ is a λ_c -value.*

Lemma 4.2. *Let $t \in \lambda_{cp}$ with $t \downarrow_{cp}$, then $enc(t) \downarrow_c$.*

Proof. Let $t_0 \in \lambda_{cp}$ with $t \downarrow_{cp}$, so $t_0 \rightarrow_{cp} t_1 \rightarrow_{cp} \dots \rightarrow_{cp} t_n$ where t_n is a value. We show by induction on n that $enc(t_0) \downarrow_c$. If $n = 0$ then t_0 is a value and $enc(t_0)$ must be a value, too, by Lemma 4.1. For the induction step we assume the induction hypothesis $enc(t_1) \downarrow_c$. Hence, it suffices to show $enc(t_0) \xrightarrow{*}_c enc(t_1)$. If $t_0 \rightarrow_{cp} t_1$ is a $(\beta\text{-CBV})$, (FIX) , $(\oplus\text{L})$, or $(\oplus\text{R})$ reduction, then the same reduction can be used in λ_c , and $enc(t_0) \rightarrow_c enc(t_1)$. If $t_0 \rightarrow_{cp} t_1$ by (SEL-F) or (SEL-S) , then three $(\beta\text{-CBV})$ steps are necessary in λ_c , i.e., $enc(t_0) \xrightarrow{3}_c enc(t_1)$. \square

For the other direction, i.e., for proving the claim $enc(t) \downarrow_c \implies t \downarrow_{cp}$ the counter example 2.1 shows that the translation enc is not adequate and not observationally correct. Moreover, this example shows that an untyped language does in general not permit an adequate – and hence also not an observationally correct – translation into a subset of itself.

$(\dots) :: \forall \alpha, \beta. \alpha \rightarrow \beta \rightarrow (\alpha, \beta)$	unit $:: \text{unit}$
fst $:: \forall \alpha, \beta. (\alpha, \beta) \rightarrow \alpha$	$\oplus :: \forall \alpha. \alpha \rightarrow \alpha \rightarrow \alpha$
snd $:: \forall \alpha, \beta. (\alpha, \beta) \rightarrow \beta$	fix $:: \forall \alpha, \beta. ((\alpha \rightarrow \beta) \rightarrow (\alpha \rightarrow \beta)) \rightarrow (\alpha \rightarrow \beta)$

Fig. 5 Types of constants

4.1 Typing λ_{cp}

One solution to prevent the counter example 2.1 is to consider a simply typed variant λ_{cp}^T of λ_{cp} as follows. The types are given by $\tau ::= \text{unit} \mid \tau \rightarrow \tau \mid (\tau, \tau)$, and only typed expressions and typed contexts are in the language λ_{cp}^T , where we assume a hole $[\cdot]_\tau$ for every type τ . For typing, we treat pairs, projections, the unit value, and the operators \oplus and **fix** as a family of constants with the types given in Fig. 5. Type safety can be stated by a preservation theorem for all expressions and a progress theorem for closed expressions. The framework now permits to prove adequacy via observational correctness of the translations.

Proposition 4.3. *For λ_{cp}^T , the (correspondingly restricted) translation $enc : \lambda_{cp}^T \rightarrow \lambda_c$ is compositional and convergence equivalent, and hence adequate.*

Proof. Compositionality follows from the definition of enc (see Fig. 4). Lemma 4.1 also holds if enc is restricted to λ_{cp}^T . We split the proof into four parts:

1. $t \downarrow_{cp} \implies enc(t) \downarrow_c$: Follows from Lemma 4.2.
2. $enc(t) \downarrow_c \implies t \downarrow_{cp}$: An inspection of the reductions shows that if t_1 is reducible, then for every reduction Red of $enc(t_1)$ to a value, there is some t_2 with $t_1 \rightarrow_{cp} t_2$ and $enc(t_1) \xrightarrow{+}_c enc(t_2)$ is a prefix of Red . We use induction on the length of a reduction Red of $enc(t)$ to a value to show that a corresponding reduction can be constructed. The base case is proved in Lemma 4.1. If t is an irreducible non-value, then due to typing it is an open expression of one of the forms $\mathbb{E}[(x r)]$, $\mathbb{E}[\mathbf{fix} x]$, $\mathbb{E}[\mathbf{fst} x]$, $\mathbb{E}[\mathbf{snd} x]$, where x is a free variable. But the cases are not possible, since $enc(t)$ is either an irreducible non-value, or $enc(t)$ reduces in one step to an irreducible non-value.
3. $enc(t) \downarrow_c \implies t \downarrow_{cp}$: We prove that $t \uparrow_{cp} \implies enc(t) \uparrow_c$ by induction on the length of a reduction $t \xrightarrow{*}_{cp} t'$, where $t' \uparrow_{cp}$. For the base case $t \uparrow_{cp}$ and (2) show that $enc(t) \uparrow_c$. The induction consists in computing a reduction sequence $enc(t) \xrightarrow{*}_c r$ where $r \uparrow_{cp}$ and the correspondence is as in the proof of Lemma 4.2, such that $t \xrightarrow{*}_{cp} t'$ and $r = enc(t')$. By type preservation, t' is well-typed and now the base-case reasoning applies.
4. $t \downarrow_{cp} \implies enc(t) \downarrow_c$: Proving $enc(t) \uparrow_c \implies t \uparrow_{cp}$ can be done using the same technique as in the previous parts. \square

Note that Proposition 3.9 cannot be applied since λ_{cp}^T is not an extension of untyped λ_c . As expected, full abstraction does not hold. For instance, let $s = \lambda p. ((\lambda y. \lambda z. (y, z)) (\mathbf{fst} p) (\mathbf{snd} p))$, and $t = \lambda p. p$. Then the equation

$$\begin{array}{l}
w \in Val_{cpig} \text{ iff } w \in Val_{cpg} \text{ iff } w \in Val_{cp} \\
t \in Exp_{cp(i)g} ::= w \mid t_1 \ t_2 \mid t_1 \oplus t_2 \mid (t_1, t_2)
\end{array}
\quad
\mathbb{E}_{cpig} ::= [] \mid \mathbb{E}_{cpig} \ t \mid w \mathbb{E}_{cpig} \mid \mid \mathbb{E}_{cpig} \oplus t \\
\mid t \oplus \mathbb{E}_{cpig} \mid (\mathbb{E}_{cpig}, t) \mid (t, \mathbb{E}_{cpig})$$

Fig. 6 Syntax of λ_{cpg} and λ_{cpig} **Fig. 7** Evaluation Contexts \mathbb{E}_{cpig} for λ_{cpig}

$$\mathbb{E}_{cpg} ::= [] \mid \mathbb{E}_{cpg} \ t \mid w \mathbb{E}_{cpg} \mid \mathbb{E}_{cpg} \oplus t \mid w \oplus \mathbb{E}_{cpg} \mid (\mathbb{E}_{cpg}, t) \mid (w, \mathbb{E}_{cpg})$$

Fig. 8 Evaluation Contexts for λ_{cpg}

$$\begin{array}{l}
enc_i : \lambda_{cpg} \rightarrow \lambda_{cp}^T : enc_i(t) = t \\
enc_g : \lambda_{cpig} \rightarrow \lambda_{cpg} : enc_g((t_1, t_2)) = (\lambda x y.(x, y)) enc_g(t_1) enc_g(t_2) \\
 : enc_g(t) = \text{descending, not changing the structure otherwise;} \\
enc_{ig} : \lambda_{cpig} \rightarrow \lambda_{cp}^T : enc_{ig} = enc_g \circ enc_i
\end{array}$$

Fig. 9 Translations between λ_{cpig} , λ_{cpg} and λ_{cp}^T

$s \sim_{cp, (\text{unit}, \text{unit}) \rightarrow (\text{unit}, \text{unit})} t$ holds in λ_{cp}^T by standard reasoning, but after translation to λ_c , we have $enc(s) \not\sim_c enc(t)$. The latter can be seen with the context $C = ([\cdot] \ \mathbf{unit})$, since $C[enc(s)]$ is must-divergent while $C[enc(t)]$ must-converges.

The extension situation could perhaps be regained by a System F-like type system, which we leave for future research. Here we just observe that the use of a simple type system for λ_c is insufficient since the encoding of pairs with components of different types cannot be simply typed. The same holds for Hindley-Milner polymorphic typing: to see this, let $s, r \in \lambda_{cp}$ where s is defined as before and $r = s \ (\mathbf{unit}, \lambda x.x)$. The most general type of $enc(s)$ in a Hindley-Milner system is $((\alpha \rightarrow \alpha \rightarrow \alpha) \rightarrow \beta) \rightarrow (\beta \rightarrow \beta \rightarrow \gamma) \rightarrow \gamma$, which essentially means that the encoding requires the components of a pair to have equal type. The reason for the insufficient type is the monomorphic use of the argument variable p of $enc(s)$. Hence, $enc(r)$ is not typeable using a Hindley-Milner type system.

One can establish a fully-abstract translation between λ_{cp}^T and a variant of λ_c by using a ‘virtual typing’ in λ_c which, intuitively, restricts λ_c to the image of the translation (see Appendix 5).

4.2 Modifying Reduction Strategies

As a final example we extend λ_{cp}^T in two steps. First, in λ_{cpg} , we allow pairs with arbitrary expressions as components (see Fig. 6). Second, in λ_{cpig} , we relax the reduction strategy by allowing interleaving evaluation of pair components and of the arguments of the choice-operator. The corresponding evaluation contexts \mathbb{E}_{cpig} for the calculus λ_{cpig} are in Fig. 7.

4.2.1 Permitting General Pairs

We consider the extension λ_{cpg} of the language λ_{cp}^T where λ_{cpg} is simply typed, and where pairs are not restricted to values. The syntax is shown in Fig. 6, the evaluation contexts in λ_{cpg} are introduced in Fig. 8. The reductions are as in λ_{cp}^T . We show that $enc_g : \lambda_{cpg} \rightarrow \lambda_{cp}^T$ is a fully abstract translation and hence nothing is lost by restricting pairs to values. Type preservation and progress also hold for λ_{cpg} . Moreover, enc_g is compositional and is easily seen to map well-typed terms of λ_{cpg} to well-typed terms of λ_{cp}^T .

Lemma 4.4. *For the translation enc_g the following holds: For all s , if s is a λ_{cpg} -value, then $enc_g(s)$ is must-convergent and has a deterministic reduction to a value. Moreover, for all s , if $enc_g(s)$ is a value, then s is a λ_{cpg} -value.*

Proof. By induction on the size of expressions and inspection of all cases. This holds also for the case $(w_1, w_2) \mapsto (\lambda x y. (x, y)) enc_g(w_1) enc_g(w_2)$, since $enc_g(w_1)$, $enc_g(w_2)$ are must-convergent and independently reduce to values, and then two deterministic beta-reductions reduce the resulting expression to a value. \square

The correctness of β -CBV-reduction in λ_{cp}^T , i.e., that $t \rightarrow t'$ by β -CBV implies $t \sim t'$, is required in the proof of full abstractness of the translation enc_g . The reason is that the composition $enc_g(-) \circ \iota$ is not the identity. In particular, pairs (x, y) of variables are not translated identically, but as $(\lambda xy. (x, y))xy$. This forces us to use Proposition 3.10, rather than Proposition 3.9.

We also require a proof of correctness of β -CBV. As a preparation, the iu-theorem from [?] is required. In contrast to the ciu-theorems found in the literature, the iu-theorem has a precondition for all value-instantiations of uses of an expression, rather than just closed instances. In the following, a *value-substitution* is defined as a substitution that replaces variables by (possibly open) values only.

Proposition 4.5 (iu-Theorem). *The following context lemma holds in λ_c . Let s, t be two λ_c -expressions of the same type τ . Then:*

- *If for all evaluation contexts E_{cp} and all value-substitutions σ : $\mathbb{E}[\sigma(s)] \Downarrow \implies \mathbb{E}[\sigma(t)] \Downarrow$, then $s \leq_1 t$ holds.*
- *If $s \leq_1 t$ and for all evaluation contexts \mathbb{E} and all value-substitutions σ : $\mathbb{E}[\sigma(s)] \Downarrow \implies \mathbb{E}[\sigma(t)] \Downarrow$, then $s \leq_{\Downarrow} t$ holds.*

Proof (sketch). This is the iu-Theorem in [?]. Before it can be applied, we have to argue that the fixpoint-reduction rule FIX can also be used in the call-by-value calculi considered in [?]. The FIX-reduction is $\mathbb{E}[\mathbf{fix} \lambda x. t] \rightarrow \mathbb{E}[t[(\lambda y. (\mathbf{fix} \lambda x. t)y)/x]]$. In order to satisfy the reduction assumptions in [?] we have to view the fixpoint-reduction as a sequence of two smaller reduction steps: $\mathbb{E}[\mathbf{fix} \lambda x. t] \rightarrow \mathbb{E}[(\lambda x. t) (\lambda y. (\mathbf{fix} \lambda x. t)y)] \rightarrow \mathbb{E}[t[\lambda y. (\mathbf{fix} \lambda x. t)y)/x]]$. Now it can be verified that the reduction assumption in [?] holds, and hence that all the proofs can be transferred. \square

Lemma 4.6. *The reduction β -CBV is a correct transformation in λ_c and also in λ_{cp}^T .*

Proof. This can be proved using the iu-theorem 4.5. Let $s = (\lambda x.r) t$ and $s' = r[t/x]$. We show that the preconditions of Proposition 4.5 hold. Since $s \rightarrow s'$, the reduction $\mathbb{E}[\sigma(s)] \rightarrow \mathbb{E}[\sigma(s')]$ holds in any case, hence $\mathbb{E}[\sigma(s')] \downarrow \implies \mathbb{E}[\sigma(s)] \downarrow$ as well as $\mathbb{E}[\sigma(s')] \Downarrow \implies \mathbb{E}[\sigma(s)] \Downarrow$. The reverse implications $\mathbb{E}[\sigma(s)] \downarrow \implies \mathbb{E}[\sigma(s')] \downarrow$ as well as $\mathbb{E}[\sigma(s)] \Downarrow \implies \mathbb{E}[\sigma(s')] \Downarrow$ also hold, since the reduction $\mathbb{E}[\sigma(s)] \rightarrow \mathbb{E}[\sigma(s')]$ is deterministic. Proposition 4.5 then implies that $s \sim s'$, and hence β -CBV is correct in λ_c . Correctness in λ_c could be proved along the same lines, or by using adequacy of the encoding $enc : \lambda_{cp}^T \rightarrow \lambda_c$ proved in Proposition 4.3. \square

Proposition 4.7. *The translation enc_g is fully abstract*

Proof. By Proposition 3.10, and since the identity: $\lambda_{cp}^T \rightarrow \lambda_{cpg}$ is an embedding (see Definition 3.8), it suffices to prove observational correctness of the translation. Note that $enc_g(t) \sim t$ due to Lemma 4.6 for all λ_{cp}^T -terms t , which makes Proposition 3.10 applicable. We have to show four implications.

1. $t \downarrow_{cpg} \implies enc_g(t) \downarrow_{cp}$: This follows by a straightforward translation from the $t \downarrow_{cpg}$ -reduction into a reduction of $enc_g(t)$. In the case of non-value pairs, (β -CBV)-reductions have to be added to produce pairs in λ_{cp}^T .
2. $enc_g(t) \downarrow_{cp} \implies t \downarrow_{cpg}$: A reduction $enc_g(t) \downarrow_{cp}$ can be re-translated into one of t , by observing that (t_1, t_2) on the λ_{cpg} -side may correspond to three different possibilities on the λ_{cp} -side: it may be (t'_1, t'_2) , $(\lambda xy.(x,y)) t'_1 t'_2$ or $(\lambda y.(t'_1, y)) t'_2$.
3. $t \Downarrow_{cpg} \implies enc_g(t) \Downarrow_{cp}$: We show $enc_g(t) \uparrow_{cp} \implies t \uparrow_{cpg}$. Again the reductions correspond, up to the (β -CBV)-reductions for the pair-encoding. The base case is that $enc_g(t) \uparrow_{cp} \implies t \uparrow_{cpg}$, which follows from (1).
4. $enc_g(t) \Downarrow_{cp} \implies t \Downarrow_{cpg}$: We show $t \uparrow_{cpg} \implies enc_g(t) \uparrow_{cp}$. As above, the reductions correspond up to the (β -CBV)-reductions for the pair-encoding. The base case is $t \uparrow_{cpg} \implies enc_g(t) \uparrow_{cp}$, and follows from (2). \square

Remark 4.8. Our version of the translation $enc_g : \lambda_{cpig} \rightarrow \lambda_{cpg}$ in a previous version of this paper was:

$$\begin{aligned} enc_{g,prev}((t_1, t_2)) &= (\lambda x y.(x,y)) enc_{g,prev}(t_1) enc_{g,prev}(t_2) && \text{if } \{t_1, t_2\} \not\subseteq Var \\ enc_{g,prev}((x,y)) &= (x,y) && \text{otherwise} \end{aligned}$$

This, however, is not compositional, since the translation is not local and depends on the surrounding context. A witness for this is the context $([\] , y)$, which cannot be translated unambiguously.

Nevertheless, this translation is compositional modulo \sim , as defined in Remark 3.11, and which can be proved using correctness of β -CBV (see Lemma 4.6). E.g. for $C = ([\] , y)$, we have $enc_{g,prev}(C[x]) = enc_{g,prev}((x,y)) = (x,y)$, $enc_{g,prev}(C[\]) = (\lambda x y.(x,y)) [\] y$, and $enc_{g,prev}(x) = x$. Now we have $enc_{g,prev}(C[x]) \sim enc_{g,prev}(C[\])(enc_{g,prev}(x))$, since β -CBV is correct w.r.t. \sim .

Remark 4.9. The combined translation from λ_{cpg} to λ_c is $enc_{gc} := enc \circ enc_g$. It operates on pairs of non-variables s, t as follows: $enc_{gc}((s,t)) = enc(\lambda xy.(x,y)) enc_{gc}(s) enc_{gc}(t) = (\lambda xy.(\lambda p.p x y)) enc_{gc}(s) enc_{gc}(t)$. The naive

translation $T'((s,t)) = (\lambda p.p T'(s) T'(t))$ is not convergence equivalent, since for example $T'((\Omega,\Omega)) = \lambda p.p \Omega \Omega$. However, (Ω,Ω) must-diverges, whereas $\lambda p.p \Omega \Omega$ is a value and thus converges.

4.2.2 Permitting Interleaved Reductions

In this subsection we will show that it is also correct to modify the reduction strategy in the OSP-calculus λ_{cpg} , where we allow that the arguments of choice and of pairs may be evaluated independently (i.e. interleaved, in any order). The OSP-calculus λ_{cpig} , i.e. its syntax and the evaluation contexts \mathbb{E}_{cpig} used for reduction have been introduced in Fig. 6 and Fig. 7. The translation $enc_i : \lambda_{cpig} \rightarrow \lambda_{cpg}$ is just the identity (see Fig. 9). However, it is not immediately obvious that the convergence predicates of λ_{cpig} and λ_{cpg} are the same, due to the independent reduction possibilities in λ_{cpig} . We denote the reduction in λ_{cpig} with \rightarrow_{cpig} and the reduction in λ_{cpg} with \rightarrow_{cpg} .

Proposition 4.10. *The identity translation enc_i from λ_{cpig} into λ_{cpg} is fully abstract.*

Proof. Obviously enc_i (and its inverse) are compositional. Thus, to prove observational correctness it suffices to establish convergence equivalence. We have to show four implications:

1. $enc_i(t) \downarrow_{cpg} \implies t \downarrow_{cpig}$: This follows by using the same reduction sequence.
2. $t \downarrow_{cpig} \implies enc_i(t) \downarrow_{cpg}$: A reduction corresponding to $t \downarrow_{cpig}$ can be rearranged until it is a reduction w.r.t. λ_{cpg} , since the reductions are at independent positions, and the final result is a value without any reductions.
3. $enc_i(t) \Downarrow_{cpg} \implies t \Downarrow_{cpig}$: We show the equivalent $t \Uparrow_{cpig} \implies enc_i(t) \Uparrow_{cpg}$. Let Red be a λ_{cpig} -reduction of $enc_i(t)$ to a must-divergent expression. We use induction on the measure (l,n) , where l is the number of reductions and n is the number of non-value surface positions of $enc_i(t)$, i.e. positions not within abstractions. Now consider the λ_{cpg} -redex in $enc_i(t)$. If the reduction of the redex is contained in Red , then we can shift it to the start, and we obtain a shorter reduction, i.e. l is decreased. Otherwise, if the reduction of the redex is not contained in Red , there are two possibilities. If the redex is must-divergent, then we are finished, since then $enc_i(t)$ is also must-divergent. Otherwise, if the redex is not must-divergent, then we simply select a converging reduction of the redex to a value. This reduction can be integrated into Red . In this case the number of reductions does not change, but the number n of the measure will be reduced. In any case, we can use induction. The base case follows from (1).
4. $t \Downarrow_{cpig} \implies enc_i(t) \Downarrow_{cpg}$: We show $enc_i(t) \Uparrow_{cpg} \implies t \Uparrow_{cpig}$. We can leave the reduction unchanged. The base case is $enc_i(t) \Uparrow_{cpg} \implies t \Uparrow_{cpig}$, which follows from (2).

Finally, full abstraction follows from Proposition 3.9, since the proof also shows that the inverse of enc_i is convergence equivalent. \square

Remark 4.11. Note that in languages with shared variable concurrency (for instance, extensions of λ_{cp} with reference cells) the modification of the reduction strategy

given in this subsection is no longer correct: permitting interleaving reductions of the arguments can be observed through their read and write effects on shared variables.

Using Proposition 3.7 we have:

Theorem 4.12. *The translation enc_{ig} is fully abstract. For $enc : \lambda_{cp}^T \rightarrow \lambda_c$ the combined translation $enc \circ enc_{ig} : \lambda_{cpig} \rightarrow \lambda_c$ is adequate.*

Conclusions and Outlook

Motivated by translation problems between concurrent programming languages, this paper succeeded in clarifying the methods, and providing tools, to assess the correctness of translations. The framework is general enough to apply directly to an operational semantics and the derived contextual equivalences, without relying on the availability of models.

In future research we want to exploit these results, to prove the correctness of various implementations of synchronization constructs in concurrent languages.

References

1. *The Alice Project*. Saarland University, <http://www.ps.uni-sb.de/alice>, 2007.
2. A. Carayol, D. Hirschhoff, and D. Sangiorgi. On the representation of McCarthy’s amb in the pi-calculus. *Theoret. Comput. Sci.*, 330(3):439–473, 2005.
3. R. De Nicola and M. Hennessy. Testing equivalences for processes. *Theoret. Comput. Sci.*, 34:83–133, 1984.
4. M. Felleisen. On the expressive power of programming languages. *Sci. Comput. Programming*, 17(1–3):35–75, 1991.
5. A. D. Gordon. Bisimilarity as a theory of functional programming. *Theoret. Comput. Sci.*, 228(1–2):5–47, 1999.
6. A. Kutzner and M. Schmidt-Schauß. A nondeterministic call-by-need lambda calculus. In *Proc. ICFP*, pages 324–335. ACM, 1998.
7. J. Matthews and R. B. Findler. Operational semantics for multi-language programs. In *34th ACM POPL*, pages 3–10. ACM, 2007.
8. G. McCusker. Full abstraction by translation. In *Advances in Theory and Formal Methods of Computing*. IC Press, 1996.
9. Robin Milner. Fully abstract models of typed lambda calculi. *Theoret. Comput. Sci.*, 4(1):1–22, 1977.
10. J. C. Mitchell. On abstraction and the expressive power of programming languages. *Sci. Comput. Programming*, 21(2):141–163, 1993.
11. J. Niehren, D. Sabel, M. Schmidt-Schauß, and J. Schwinghammer. Observational semantics for a concurrent lambda calculus with reference cells and futures. *Electron. Notes Theor. Comput. Sci.*, 173:313–337, 2007.
12. J. Niehren, J. Schwinghammer, and G. Smolka. A concurrent lambda calculus with futures. *Theoret. Comput. Sci.*, 364(3):338–356, 2006.
13. S. Peyton Jones, A. Gordon, and S. Finne. Concurrent Haskell. In *23rd ACM POPL*, pages 295–308. ACM, 1996.

14. A. D. Pitts. Parametric polymorphism and operational equivalence. *Math. Structures Comput. Sci.*, 10:321–359, 2000.
15. J. G. Riecke. Fully abstract translations between functional languages. In *18th ACM POPL*, pages 245–254. ACM, 1991.
16. E. Ritter and A. M. Pitts. A fully abstract translation between a lambda-calculus with reference types and Standard ML. In *Proc. 2nd TLCA*, pages 397–413. Springer, 1995.
17. D. Sabel and M. Schmidt-Schauß. A call-by-need lambda-calculus with locally bottom-avoiding choice: Context lemma and correctness of transformations. *Math. Structures Comput. Sci.*, 2008. accepted for publication.
18. S. B. Sanjani and C.-H. L. Ong. Fully abstract semantics of additive aspects by translation. In *Proc. 6th OASD*, pages 135–148. ACM, 2007.
19. Manfred Schmidt-Schauß and David Sabel. On generic context lemmas for lambda calculi with sharing. Frank report 27, Institut für Informatik. Fachbereich Informatik und Mathematik. J. W. Goethe-Universität Frankfurt am Main, August 2007.
20. E. Shapiro. Separating concurrent languages with categories of language embeddings. In *23rd ACM STOC*, pages 198–208. ACM, 1991.
21. H. Søndergaard and P. Sestoft. Non-determinism in functional languages. *Comput. J.*, 35(5):514–523, 1992.

5 Using Virtual Typing

We consider the issue of full abstraction of the pair encoding for simply typed λ_{cp}^T and assume a variation of simple typing for λ_c in order to describe the structure of the image of λ_{cp}^T under the translation $enc(\cdot)$.

We define the language λ_c^{VT} as a typed variant of λ_c that is sufficiently large to serve as a target language for enc . The syntax of expressions in λ_c^{VT} is extended as follows. We assume that every expression s and subexpression is decorated with a pair $\langle \tau, \beta \rangle$ of labels: a type label τ , and a *selector-label* β , written as $s :: \langle \tau, \beta \rangle$, where we write $s :: \tau$, if only the type label is of interest. Here, τ is either a λ_{cp}^T -type (i.e., including pair types), or the special symbol \dagger (indicating no type), and the selector-label can be either **fst**, **snd** or $\#$. Intuitively, $\#$ can be interpreted as the absence of a selector-label. The objects of the language λ_c^{VT} are triples $(s, \langle \tau, \beta \rangle)$. Thus there may be different objects corresponding to the same λ_c -expression. Below, we give more conditions that will only accept certain triples as valid λ_c^{VT} -expressions. We assume that variables are partitioned by assigning a fixed type (or \dagger) to each, which is also its type-label. Constants are labeled with a type that is an instance of the type as given in Fig. 5.

Instead of type derivation rules, we assume that the following consistency rules must be satisfied by λ_c^{VT} -expressions and their type- and selector-labeling. That is, types are not inferred for expressions and subexpressions, but verified against the term structure and the type- and selector-labeling.

Definition 5.1. The type consistency rules for λ_c^{VT} are:

- An application $(s t)$ is type-labeled as: $((s :: \tau_1 \rightarrow \tau_2) (t :: \tau_1)) :: \tau_2$, if it is not one of the exceptions in 2 below.
- There are different possibilities for abstractions. The expression $\lambda x.t$ is consistently typed if the term and the type labeling satisfies one of the following patterns:
 1. $\lambda(x :: \tau_1).(t :: \tau_2) :: \tau_1 \rightarrow \tau_2$.
 2. $(\lambda(sel :: \dagger).sel(w_1 :: \tau_1)(w_2 :: \tau_2)) :: (\tau_1, \tau_2)$, the variable sel does not occur free in s, t , w_1, w_2 are λ_c -values after stripping off the labels, and w_1, w_2 have to be type-consistent. Type consistency is not necessary for the applications $(sel(w_1 :: \tau_1))$ and $sel(w_1 :: \tau_1)(w_2 :: \tau_2)$. This kind of expression is the only possibility for a variable to be labeled with \dagger as a type.
 3. $(\lambda p.p(\lambda x.\lambda y.x)) :: \langle (\tau_1, \tau_2) \rightarrow \tau_1, \mathbf{fst} \rangle$ where $p :: ((\tau_1 \rightarrow \tau_2 \rightarrow \tau_1) \rightarrow \tau_1)$; $(\lambda x.\lambda y.x) : (\tau_1 \rightarrow \tau_2 \rightarrow \tau_1)$.
 4. $(\lambda p.p(\lambda x.\lambda y.y)) :: \langle (\tau_1, \tau_2) \rightarrow \tau_2, \mathbf{snd} \rangle$ where $p :: ((\tau_1 \rightarrow \tau_2 \rightarrow \tau_2) \rightarrow \tau_2)$; $(\lambda x.\lambda y.y) : (\tau_1 \rightarrow \tau_2 \rightarrow \tau_2)$.

Whenever an expression has a selector-label **fst** or **snd**, then it must be one of the cases (3),(4) above.

Typing of the constants is as for λ_{cp}^T . An expression t of λ_c^{VT} is *well-typed* of type τ if t is type-labeled τ , and the type consistency rules hold for the subexpression of

t according to Definition 5.1. A λ_c^{VT} -value is defined to be a (labeled) abstraction or a constant.

The action of the reduction rules in λ_c^{VT} on the expressions and hence the label components is the obvious one, with the exception of the cases where the redex is an application of a selector-labeled expression to a pair, which is defined explicitly:

Definition 5.2. We define the type behavior of the reduction rules in λ_c^{VT} for the critical cases of an application of an implemented selector to a pair.

- Let the redex be an application $(s t)$, where s has selector-label **fst**,
 $s = \lambda p.p (\lambda x.\lambda y.x) :: (\tau_1, \tau_2) \rightarrow \tau_1$, $t :: (\tau_1, \tau_2)$, $(s t) :: \tau_1$, the term t must be an abstraction
 $(\lambda (sel :: \dagger).sel (w_1 :: \tau_1) (w_2 :: \tau_2)) :: (\tau_1, \tau_2)$.
 Then the beta-reduction will produce the expression

$$(\lambda (sel' :: \tau_1 \rightarrow \tau_2 \rightarrow \tau_1).sel' (w_1 :: \tau_1) (w_2 :: \tau_2)) (\lambda x.\lambda y.x)$$

with the type label τ_1 . The selector-label **fst** is removed.

- Similarly for selector-label **snd**.
- Beta-reduction must give priority to the selector-labels **fst, snd** over the label **#**. The latter may be overwritten.

Note that the reductions on the underlying λ_c -expression are exactly the same as the untyped reductions in λ_c .

A case-analysis results in the following:

Proposition 5.3. *The following holds for λ_c^{VT} :*

- *The type of a closed expression is not changed by reduction.*
- *A closed well-typed expression is either reducible or a value.*

As observers in λ_c^{VT} we use the contexts of λ_c^{VT} with the following restrictions: The hole is also typed, an expression with a selector-label cannot have a hole in it; and a context or a term cannot contain a free variable with type label \dagger .

The translations $enc(\cdot)$ and $inc(\cdot)$ are adapted to the labeling (see Fig. 10 and 11): $enc(\cdot)$ keeps the type labeling and adds the select-labels. The translation $inc(\cdot)$ maps abstractions to pairs controlled by the type labeling, and uses the selector-labels to map abstractions to the appropriate selectors.

The type labeling of contexts shows that the translations between λ_{cp}^T and λ_c^{VT} are compositional, and that the type mapping is the identity.

It is easy to verify that if s is a well-typed expression in λ_{cp}^T , then $enc(s)$ is well-typed as a λ_c^{VT} -expression, and conversely if s in λ_c^{VT} is well-typed then $inc(s)$ is well-typed as a λ_{cp}^T -expression. Thus enc and inc are translations also w.r.t. the typed languages.

Lemma 5.4. *For the mappings enc and inc , the following holds:*

- *For all s : s is a λ_{cp}^T -value iff $enc(s)$ is a value.*
- *For all s : s is a λ_c^{VT} -value iff $inc(s)$ is a value.*

$$\begin{aligned}
\mathit{enc}((w_1, w_2) :: (\tau_1, \tau_2)) &= \lambda s. ((s :: \dagger) \mathit{enc}(w_1) :: \tau_1 \mathit{enc}(w_2) :: \tau_2) :: (\tau_1, \tau_2) \\
\mathit{enc}(\mathbf{fst} :: (\tau_1, \tau_2) \rightarrow \tau_1) &= \lambda p. (p \lambda x. \lambda y. x) :: ((\tau_1, \tau_2) \rightarrow \tau_1, \mathbf{fst}) \\
\mathit{enc}(\mathbf{snd} :: (\tau_1, \tau_2) \rightarrow \tau_2) &= \lambda p. (p \lambda x. \lambda y. y) :: ((\tau_1, \tau_2) \rightarrow \tau_2, \mathbf{snd})
\end{aligned}$$

Fig. 10 Adaptations of translation enc

$$\begin{aligned}
\mathit{inc}(\lambda s. ((s :: \dagger) w_1 :: \tau_1 w_2 :: \tau_2) :: ((\tau_1, \tau_2), \#)) &= (\mathit{inc}(w_1), \mathit{inc}(w_2)) :: (\tau_1, \tau_2) \\
\mathit{inc}(\lambda p. (p \lambda x. \lambda y. x) :: ((\tau_1, \tau_2) \rightarrow \tau_1, \mathbf{fst})) &= \mathbf{fst} :: (\tau_1, \tau_2) \rightarrow \tau_1 \\
\mathit{inc}(\lambda p. (p \lambda x. \lambda y. y) :: ((\tau_1, \tau_2) \rightarrow \tau_2, \mathbf{snd})) &= \mathbf{snd} :: (\tau_1, \tau_2) \rightarrow \tau_2
\end{aligned}$$

Only the top-types are indicated and the \dagger -label of s .

Fig. 11 Adaptations of translation inc

Proof. Follows by inspecting all the cases. \square

Theorem 5.5. *The translations inc and enc are fully abstract translations between λ_{cp}^T and λ_c^{VT} .*

Proof. Compositionality follows from the definition of the translations. Values are preserved and observational correctness of inc holds. The proof of observational correctness of enc is analogous to the proof of Proposition 4.3.

Full abstraction of both translations holds, since types are preserved, and the translations are partial inverses of each other:

$\mathit{enc}(\mathit{inc}(s)) = s$ for each λ_c^{VT} -expression s (modulo α -renaming), and

$\mathit{inc}(\mathit{enc}(s)) = s$ for each λ_{cp}^T -expression s (modulo α -renaming).

Then it is easy to show that the observers are also equivalent, and hence that full abstraction of both translations holds. \square

This result can be interpreted as an isomorphism between λ_{cp}^T and λ_c^{VT} , which may sloppily be formulated as “the pair-constructor can be encoded in the types”. However, note that λ_{cp}^T is not an extension of λ_c^{VT} , since the λ_c^{VT} -expressions are not contained in λ_{cp}^T . In particular, an abstraction in λ_{cp}^T may correspond to several objects in λ_c^{VT} due to the type labeling.