Correctly Translating Concurrency Primitives

Jan Schwinghammer

Saarland University, Saarbrücken,
Germany

David Sabel

Manfred Schmidt-Schauß
Goethe-University, Frankfurt, Germany

Joachim Niehren

INRIA, Lille, France, Mostrare Project

Abstract

Motivated by the question of correctness of a specific implementation of concurrent buffers in the lambda calculus with futures underlying Alice ML, we prove that concurrent buffers and handled futures can correctly encode each other. Our translations map waiting on handled futures to queuing of concurrent buffers and vice versa. Correctness of translations means that they preserve and reflect the observations of may- and must-convergence. As a consequence of compositionality, they are also adequate with respect to a contextually defined notion of observational program semantics. We demonstrate that our approach to the correctness of implementations applies uniformly to the whole compilation process from high-level to low-level concurrent languages.

Categories and Subject Descriptors D.3.3 [Programming Languages]: Language Constructs and Features – *concurrent programming structures*; F.3.1 [Logics and Meanings of Programs]: Specifying and Verifying and Reasoning about Programs; F.3.2 [Logics and Meanings of Programs]: Semantics of Programming Languages – *operational semantics*

General Terms Theory, verification

Keywords Concurrency, lambda calculus, semantics

1. Introduction

Modern concurrent programming languages extend sequential languages with concurrent threads and concurrency primitives for controlling their interactions. Computation within each thread is sequential. Examples for concurrency primitives are MVars (i.e. concurrent buffers) in Haskell (Jones et al. 1996), channels in Concurrent ML (Reppy 1999), handled futures in Alice ML (Rossberg et al. 2006), and joins in JoCaml (Fournet et al. 2002).

Alice ML is a concurrent extension of Standard ML, adding both eager and lazy threads. The objective of the present paper is to prove precise relationships between synchronization primitives in the context of Alice ML, in particular the relation between handled futures and concurrent buffers. Our first motivation is to clarify the choice of the language primitives in the design of the Alice ML language. The second motivation is to prove the correctness of the implementation of concurrent buffers in Alice ML. We use the *concurrent lambda calculus with futures* (Niehren

et al. 2006) to model the operational semantics of the concurrent core of Alice ML. Correctness results are expressed with respect to the observational semantics for this calculus, that is based on may- and must-convergence. The third motivation is to demonstrate the usefulness of recent results and proof techniques for observational semantics for concurrent languages.

The main contribution of this paper is a comprehensive proof that concurrent buffers and handled futures are equivalent synchronization primitives. From a semantic point of view, the designers of Alice ML can thus choose either concept as a primitive and the other as a library contribution. Both translations are compositional and preserve may- and must-convergence, which means that they are ‘observationally correct’ and allows us to transfer program equivalences back and forth. Thus, the recently developed proof techniques based on observational semantics for concurrent languages (Schmidt-Schauß et al. 2008) have been usefully applied in an ambitious setting.

Language translations and observational semantics. We now describe our approach and results in more detail. We start with an enriched core language of Alice ML, the calculus $\lambda^\tau(\text{fch})$ which extends the calculi of Niehren et al. (2006, 2007). This is a typed call-by-value lambda with futures, polymorphic data and type constructors, concurrent threads, reference cells, and handled futures. Programs in this language consist of a collection of eager and lazy threads. When a concurrent thread is spawned, it immediately returns a *future*, which is a placeholder for the value computed by this thread. Other threads can proceed with this placeholder until the actual value is needed, in which case they block on the future, and they resume once the value becomes available. Besides these concurrent futures associated with threads, there are futures with an explicit *handler* which supports single assignment of values to futures. Handled futures, one of the two synchronization primitives we study here, are called *promises* in Alice ML, following the proposal of Liskov and Shriram (1988).

We use a contextual observational semantics for the concurrent lambda calculus with futures, based on operationally-defined forms of may- and must-convergence (De Nicola and Hennessy 1984; Ong 1993; Carayol et al. 2005). Our form of must-convergence is similar to the should-testing of Rensink and Vogler (2007). A common feature is that fairness of execution is mirrored in the semantic theory. This combination of may- and must-convergence properly captures the non-determinism arising in concurrent programming languages (Sabel and Schmidt-Schauß 2008; Niehren et al. 2007). Given a language $\mathcal{C}$ we write $=_{\mathcal{C}}$ for the observational semantics on programs of the language $\mathcal{C}$, which equates all programs with equal may- and must- convergence behavior in all contexts.

For correctness of translations between languages $T : \mathcal{C} \rightarrow \mathcal{C}'$, we use the notion of *observational correctness*, which for compositional translations means that all programs p and $T(p)$ exhibit the same convergence behavior. Observational correctness implies *adequacy* with respect to observational semantics $=_{\mathcal{C}}$ and $=_{\mathcal{C}'}$ (Riecke

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.

ML'09, August 30, 2009, Edinburgh, Scotland, UK.

Copyright © 2009 ACM 978-1-60558-509-3/09/08...\$5.00

1991; Ritter and Pitts 1995; Schmidt-Schauß et al. 2008), meaning that all program transformations of $\mathcal{C}'$ can be soundly applied on $\mathcal{C}$ up to the translation T . Formally, a translation T is adequate if all (equally typed) programs with equivalent translations are equivalent, i.e., if $T(p_1) =_{\mathcal{C}'} T(p_2)$ implies $p_1 =_{\mathcal{C}} p_2$. If additionally the converse holds, then T is called fully abstract. While full abstraction follows from observational correctness and some additional assumptions, it plays only a minor role for our results.

Proving observational correctness of translations. In previous work (Niehren et al. 2007) we analyzed a less expressive untyped core language, the calculus $\lambda(\text{fh})$. In comparison to the calculus $\lambda^\tau(\text{fch})$ considered here, it lacks data constructors and case expressions which are critical for our specification of buffers. For this core language we have proved a rich set of program transformations correct with respect to contextual equivalence, using diagram-based techniques based on the operational semantics. Instead of applying this technically involved and complex mechanism again to $\lambda^\tau(\text{fch})$, we use adequacy to lift the correctness results obtained for $\lambda(\text{fh})$ to $\lambda^\tau(\text{fch})$, by finding suitable adequate translations.

Equipped with these results about the equational theory we focus on the correctness of implementations of buffers in $\lambda^\tau(\text{fch})$. To obtain a specification of buffers that is sufficiently rigorous for a correctness argument we extend $\lambda^\tau(\text{fch})$ by concurrent buffers (as e.g. used for implementing buffered channels by Peyton Jones et al. (1996)), resulting in the calculus $\lambda^\tau(\text{fchb})$. We then present an implementation of buffers into the buffer-free calculus, such that blocking buffer operations are translated into queuing and waiting. After formalizing this implementation as a translation we show its adequacy. Moreover, the translation turns out to be observationally correct. As mentioned above, this means that the specification and implementation of buffers give rise to the same observations (in particular, they have the same convergence behavior).

We complement our result by showing that it is also possible to go into the opposite direction and correctly implement handled futures with buffers. In this case, the specification is again the calculus $\lambda^\tau(\text{fchb})$, but now viewed as an extension of a handle-free calculus with buffers, called $\lambda^\tau(\text{fcb})$. We provide a translation from $\lambda^\tau(\text{fchb})$ to $\lambda^\tau(\text{fcb})$ and show that it is fully abstract.

The implementations of buffers and handles lead to the question whether the constructs can already be encoded in a base language containing neither buffers nor handles. We show that this is indeed the case. However, the implementation that underlies this last translation can result in busy-wait situations, indicating that it is of a more low-level character than the other encodings.

For proving adequacy, full abstraction and observational correctness of the various encodings we rely on compositionality (Schmidt-Schauß and Sabel 2007; Schmidt-Schauß et al. 2008), on commutation methods in order to prove invariants of implicit queuing mechanisms that arise when implementing buffers by handled futures, as well as on equivalences for $\lambda^\tau(\text{fch})$ that we lift from $\lambda(\text{fh})$. In turn, we inherit equations for $\lambda^\tau(\text{fchb})$ from $\lambda^\tau(\text{fch})$ by the adequacy of T_B .

Summary and results. Table 1 presents the syntactic features of the various calculi which are considered in this paper.

calculus	typed	data constructors	handled futures	buffer primitives
$\lambda(\text{fh})$	×	×	✓	×
$\lambda^\tau(\text{fc})$	✓	✓	×	×
$\lambda^\tau(\text{fch})$	✓	✓	✓	×
$\lambda^\tau(\text{fcb})$	✓	✓	×	✓
$\lambda^\tau(\text{fchb})$	✓	✓	✓	✓

Table 1. Overview of the calculi considered in this paper

Our main interest lies in the last three calculi, $\lambda^\tau(\text{fch})$, $\lambda^\tau(\text{fcb})$, and $\lambda^\tau(\text{fchb})$, within which we formulate the implementations of buffers and of futures. The following diagram summarizes our results. Doubly lined arrows ($\Longrightarrow$) indicate fully abstract translations, while single lined arrows ($\longrightarrow$) indicate adequate translations.

The translation T_B describes the implementation of concurrent buffers in terms of handled futures. Conversely, the translation T_H realizes the implementation of handled futures with buffers. The translations ι_B and ι_H represent the embeddings of the smaller calculi into the extended calculus $\lambda^\tau(\text{fchb})$ that contains both buffers and handled futures as primitives. All the translations shown in the diagram are observationally correct, and thus also preserve the convergence behavior, i.e., all these implementations are indistinguishable from the respective primitive constructs. Note that we left the exact status of T_B open: we do not know if it is fully abstract.

Of course there are additional encodings which are implied by the ones shown in the diagram. In particular, we mention that we obtain adequate encodings between $\lambda^\tau(\text{fcb})$ into $\lambda^\tau(\text{fch})$.

Outline. The next section presents the calculus $\lambda^\tau(\text{fch})$ underlying Alice ML, and also abstracts from its observational semantics a more general framework with respect to which we can define language translations and establish some of their basic properties. Section 3 develops the equational theory of $\lambda^\tau(\text{fch})$, by showing the adequacy of translation to $\lambda(\text{fh})$. This paves the way for the correctness proof for the implementation of buffers, given in Section 4. Section 5 considers the implementation of handled futures by buffers (T_H in the diagram above), and Section 6 considers an implementation of buffers in a calculus without handled futures (T_{RB} above), based on busy waiting. Open questions and related work are discussed in Section 7.

More explanations and full proofs can be found in the accompanying technical report (Schwinghammer et al. 2009).

2. Lambda Calculus with Futures and Constructors

This section presents the calculus $\lambda^\tau(\text{fch})$ underlying Alice ML. This is a typed lambda calculus with algebraic data types, concurrent and handled futures, and reference cells, which is obtained from the calculus with futures of Niehren et al. (2006) by adding data constructors with recursive polymorphic type constructors.

Type and data constructors. Our encodings require n -tuples $\langle v_1, \dots, v_n \rangle$ of all possible types $\tau_1 \times \dots \times \tau_n$. For the sake of generality and uniformity, we keep the concrete signature of data and type constructors as a parameter. Such a signature $\Sigma = (\mathcal{K}, \mathcal{D})$ consists of a finite ranked set of type constructors $\kappa \in \mathcal{K}$ and a finite ranked set of data constructors $k \in \mathcal{D}$. We denote the arities of data and type constructors by $ar(\cdot) \geq 0$. Polymorphic types $\hat{\tau}$ over Σ have the following abstract syntax, where α belongs to a fixed infinite set of type variables:

$$\hat{\tau} \in PolyType ::= \alpha \mid \text{unit} \mid \text{ref } \hat{\tau} \mid \hat{\tau} \rightarrow \hat{\tau} \mid \kappa(\hat{\tau}_1, \dots, \hat{\tau}_{ar(\kappa)})$$

Monomorphic types $\tau \in Type$ are polymorphic types without variables. We assume a unique polymorphic type $\text{upt}(k) \in PolyType$ for each data constructor $k \in \mathcal{D}$, that has the form $\hat{\tau}_1 \rightarrow \dots \rightarrow \hat{\tau}_{ar(k)} \rightarrow \kappa(\alpha_1, \dots, \alpha_{ar(\kappa)})$ where $\kappa \in \mathcal{K}$ and only $\alpha_1, \dots, \alpha_{ar(\kappa)}$ may occur as type variables in $\hat{\tau}_j$ for all $j =$

$$\begin{aligned}
\tau &\in \text{Type} ::= \text{unit} \mid \text{ref } \tau \mid \tau \rightarrow \tau \mid \kappa(\tau_1, \dots, \tau_{\text{ar}(\kappa)}) \\
c &\in \text{Const} ::= \text{unit} \mid \text{ref}^\tau \mid \text{thread}^\tau \mid \text{lazy}^\tau \mid \text{handle}^\tau \\
\pi &\in \text{Pat} ::= k^\tau(x_1, \dots, x_{\text{ar}(k)}) \\
e &\in \text{Exp} ::= x \mid c \mid \lambda x. e \mid e_1 e_2 \mid \text{exch}(e_1, e_2) \\
&\quad \mid k^\tau(e_1, \dots, e_{\text{ar}(k)}) \\
&\quad \mid \text{case}_\kappa e \text{ of } \pi_1 \Rightarrow e_1 \mid \dots \mid \pi_m \Rightarrow e_m \quad (m > 0) \\
v &\in \text{Val} ::= x \mid c \mid \lambda x. e \mid k^\tau(v_1, \dots, v_{\text{ar}(k)}) \\
p &\in \text{Proc} ::= p_1 \mid p_2 \mid (\nu x)p \mid x \text{ c } v \mid x \Leftarrow e \\
&\quad \mid x \xleftarrow{\text{susp}} e \mid y \text{ h } x \mid y \text{ h } \bullet
\end{aligned}$$

Figure 1. Types, expressions and processes of $\lambda^\tau(\text{fch})$

$$\begin{aligned}
p_1 \mid p_2 &\equiv p_2 \mid p_1 & (p_1 \mid p_2) \mid p_3 &\equiv p_1 \mid (p_2 \mid p_3) \\
(\nu x)(\nu y)p &\equiv (\nu y)(\nu x)p & (\nu x)(p_1) \mid p_2 &\equiv (\nu x)(p_1 \mid p_2) \\
&& \text{if } x \notin \text{fv}(p_2)
\end{aligned}$$

Figure 2. Structural congruence of processes

$1, \dots, \text{ar}(k)$. The set $\mathcal{D}(\kappa)$ consists of all data constructors $k \in \mathcal{D}$ for which κ occurs in the target type of $\text{upt}(k)$. We assume that $\mathcal{D}(\kappa)$ is nonempty for all $\kappa \in \mathcal{K}$. In typing rules, we will write $\tau \preceq \hat{\tau}$ if τ is a monomorphic instance of the polymorphic type $\hat{\tau} \in \text{PolyType}$.

For instance, we can define lists of all types, when having a type constructor $\text{List} \in \mathcal{K}$ with two data constructors $\text{cons}, \text{nil} \in \mathcal{D}(\text{List})$ such that $\text{upt}(\text{cons}) = \alpha \rightarrow \text{List}(\alpha) \rightarrow \text{List}(\alpha)$ and $\text{upt}(\text{nil}) = \text{List}(\alpha)$. For n -tuples (where $n \geq 0$), we assume type constructors $\cdot \times \dots \times \cdot \in \mathcal{K}$ and data constructors $\langle \cdot, \dots, \cdot \rangle \in \mathcal{D}$ of arity n , whose unique polymorphic type is $\alpha_1 \rightarrow \dots \rightarrow \alpha_n \rightarrow (\alpha_1 \times \dots \times \alpha_n)$.

Syntax and typing of $\lambda^\tau(\text{fch})$. We start from a signature $(\mathcal{K}, \mathcal{D})$, a set of variables Var , and a global assignment of variables to monomorphic types $\Gamma : \text{Var} \rightarrow \text{Type}$ such that for every $\tau \in \text{Type}$ there exists infinitely many $x \in \text{Var}$ with $\Gamma(x) = \tau$. Consistent renaming of variables must preserve the type. The syntax of $\lambda^\tau(\text{fch})$ defined in Fig. 1 consists of two layers: a level of λ -expressions $e \in \text{Exp}$ for sequential computation within threads, and a level of processes $p \in \text{Proc}$ that compose threads in parallel and record the state of the system. Expressions e subsume values v as usual in a call-by-value λ -calculus. Compared to the original lambda calculus with futures, we add constructor applications $k^\tau(e_1, \dots, e_{\text{ar}(k)})$ creating data structures by constructors $k \in \mathcal{D}$ labeled with monomorphic types $\tau \in \text{Type}$, and typed case-expressions $\text{case}_\kappa e \text{ of } \pi_1 \Rightarrow e_1 \mid \dots \mid \pi_m \Rightarrow e_m$ whose pattern are non-overlapping and exhaustive. Thus, every constructor $k \in \mathcal{D}(\kappa)$ appears exactly once in some pattern π_i . A pattern has the form $k^\tau(x_1, \dots, x_{\text{ar}(k)})$ such that no variable appears twice. All variables in a pattern π of a branch $\pi \Rightarrow e$ are bound with scope in e . The set of free variables of e is denoted by $\text{fv}(e)$ ($\text{fv}(p)$ for processes p). Expressions and processes are identified up to consistent renaming of bound variables.

New components in expressions are introduced by typed (higher-order) constants. The constant ref^τ introduces a reference cell. The constants thread^τ and lazy^τ serve for introducing eager threads and lazy threads, each of them together with a future. Finally, the constant handle^τ is used to generate futures with an associated handler. For convenience we sometimes omit the type label of constants as well as constructors if it is obvious or not important. The expression $\text{exch}(e_1, e_2)$ expresses atomic exchange of cell values. Note that we distinguish between constants and data constructors $k \in \mathcal{D}$ — the latter must always be fully applied.

$$\begin{aligned}
&\frac{\Gamma(x) = \tau}{x : \tau} \quad \frac{}{\text{unit} : \text{unit}} \quad \frac{\tau \preceq \alpha \rightarrow \text{ref } \alpha}{\text{ref}^\tau : \tau} \quad \frac{\tau \preceq (\alpha \rightarrow \alpha) \rightarrow \alpha}{\text{thread}^\tau : \tau} \\
&\frac{\tau \preceq (\alpha \rightarrow \alpha) \rightarrow \alpha}{\text{lazy}^\tau : \tau} \quad \frac{\tau \preceq (\alpha_1 \rightarrow (\alpha_1 \rightarrow \text{unit}) \rightarrow \alpha_2) \rightarrow \alpha_2}{\text{handle}^\tau : \tau} \\
&\frac{x : \tau_1 \quad e : \tau_2}{(\lambda x. e) : \tau_1 \rightarrow \tau_2} \quad \frac{e_1 : \tau_1 \rightarrow \tau_2 \quad e_2 : \tau_1}{(e_1 e_2) : \tau_2} \quad \frac{e_1 : \text{ref } \tau \quad e_2 : \tau}{\text{exch}(e_1, e_2) : \tau} \\
&\frac{k \in \mathcal{D}(\kappa) \quad \forall j \in 1 \dots \text{ar}(k). e_j : \tau_j \quad \tau = \tau_1 \rightarrow \dots \rightarrow \tau_{\text{ar}(k)} \rightarrow \kappa(\tau'_1, \dots, \tau'_{\text{ar}(\kappa)}) \preceq \text{upt}(k)}{k^\tau(e_1, \dots, e_{\text{ar}(k)}) : \kappa(\tau'_1, \dots, \tau'_{\text{ar}(\kappa)})}
\end{aligned}$$

$$\begin{aligned}
\mathcal{D}(\kappa) &= \{k_1, \dots, k_n\} \quad e : \kappa(\tau'_1, \dots, \tau'_{\text{ar}(\kappa)}) \quad \forall i = 1 \dots n. e_i : \tau \\
&\quad \forall i = 1 \dots n. \tau_i = \Gamma(x_{i,1}) \rightarrow \dots \rightarrow \Gamma(x_{i,\text{ar}(k_i)}) \rightarrow \kappa(\tau'_1, \dots, \tau'_{\text{ar}(\kappa)}) \preceq \text{upt}(k_i) \\
&\quad (\text{case}_\kappa e \text{ of } (k_i^{\tau_i}(x_{i,1}, \dots, x_{i,\text{ar}(k_i)}) \Rightarrow e_i)^{i=1 \dots n}) : \tau
\end{aligned}$$

Figure 3. Types of expressions

As in the pi-calculus, processes p are composed from smaller components by parallel composition $p_1 \mid p_2$ and new name creation $(\nu x)p$. The latter is a variable binder. It can be seen as hiding variables, whereas free variables are visible for outside observers. A *structural congruence* $\equiv$ on processes is defined by the axioms in Fig. 2. We distinguish five types of components that have no direct correspondence in pi-calculus. Cells $x \text{ c } v$ associate (a memory location) x to a value v . Eager concurrent threads $x \Leftarrow e$ will eventually bind future x to the value of expression e unless it diverges or suspends; x is called a *concurrent future*. Lazy threads $x \xleftarrow{\text{susp}} e$ are suspended computations that will start once the proper value of x is needed elsewhere; we call x a *lazy future*. Handle components $y \text{ h } x$ associate handles y to futures x , so that y can be used to assign a value to x . We call x a *future handled by y* , or more shortly a *handled future*. Finally, a used handle component $y \text{ h } \bullet$ indicates that y is a handle that has already been used to bind its associated future. A process p *introduces* a variable x if $p \equiv p_1$ or $p \equiv p_1 \mid p_2$ for p_1 a component of the following form (for some e, v and y): $x \text{ c } v$, or $x \Leftarrow v$, or $x \xleftarrow{\text{susp}} e$, or $y \text{ h } x$, or $x \text{ h } y$, or $x \text{ h } \bullet$. A process is *well-formed* if no subprocess introduces any variable more than once. For instance, neither $x \Leftarrow v \mid x \text{ c } v'$ nor $(\nu x)(x \Leftarrow v \mid x \text{ c } v')$ is well-formed.

In order to have a consistent notion of typed program transformation, we rely on unique monomorphic typings. To this end, we already assumed a unique type $\Gamma(x)$ for all variables. For expressions, we assign types in judgements $e : \tau$. Process components have to be well-typed, written $p : \text{wt}$. The typing rules for expressions and processes can be found in Fig. 3 and 4. Note that the well-formedness conditions for processes described earlier are kept orthogonal to typing, in contrast to the type system of Niehren et al. (2006). We write $e[e'/x]$ for the (capture-free) substitution of x by e' in e . It preserves the type of e if $e' : \Gamma(x)$ holds.

Syntactic abbreviations. We assume that the set of type constructors contains a nullary constructor $\text{bool} \in \mathcal{K}$, with nullary data constructors true and false . For convenience, we will freely use the usual syntactic sugar such as a (non-recursive) let-binding $\text{let } x_1 = e_1, \dots, x_n = e_n \text{ in } e$ and sequencing $e_1; e_2$, and also use patterns in abstractions $\lambda \pi. e$ as shorthand for $\lambda x. \text{case } e \text{ of } \pi \Rightarrow e; \pi' \Rightarrow z$ etc. (where z represents an error, and for instance can be defined by the component $z \Leftarrow z$). Instead of $\text{case } e \text{ of true} \Rightarrow e_1 \mid \text{false} \Rightarrow e_2$ we write if e then e_1 else e_2 , and the special case if e then true else true is written as *wait e* . The

$\frac{p_1 : wt \quad p_2 : wt}{p_1 \mid p_2 : wt}$	$\frac{x:\tau \quad e:\tau}{x \leftarrow e : wt}$	$\frac{x:\tau \quad e:\tau}{x \xleftarrow{susp} e : wt}$	$\frac{x:\text{ref } \tau \quad v:\tau}{x \text{ c } v : wt}$
$\frac{p:wt}{(\nu x)p : wt}$	$\frac{}{y \text{ h } \bullet : wt}$	$\frac{x:\tau \quad y:\tau \rightarrow \text{unit}}{y \text{ h } x : wt}$	

Figure 4. Well-typed processes.

ECs	$E ::= x \leftarrow \tilde{E}$ $\tilde{E} ::= []^\tau \mid \tilde{E} e \mid v \tilde{E} \mid \text{exch}(\tilde{E}, e) \mid \text{exch}(v, \tilde{E})$ $\mid \text{case } \tilde{E} \text{ of } (\pi_i \Rightarrow e_i)^{i=1 \dots n}$ $\mid k(v_1, \dots, v_{i-1}, \tilde{E}, e_{i+1}, \dots, e_n)$
Future ECs	$F ::= x \leftarrow \tilde{F}$ $\tilde{F} ::= \tilde{E}[[\cdot]^\tau v] \mid \tilde{E}[\text{exch}([\cdot]^\tau, v)]$ $\mid \tilde{E}[\text{case } [\cdot]^\tau \text{ of } (\pi_i \Rightarrow e_i)^{i=1 \dots n}]$
Process ECs	$D ::= [] \mid p \mid D \mid p \mid (\nu x)D$

Figure 5. Evaluation contexts

symbol ‘ $\cdot$ ’ stands for an arbitrary fresh variable. Finally, we write newhandled as shorthand for **handle** $\lambda f \lambda h. \langle h, f \rangle$.

Contexts and operational semantics. The operational semantics defines an evaluation strategy via evaluation contexts in which reduction rules apply. We introduce *contexts* C and D . An *expression context* C is a process where exactly one expression-position is replaced with a typed hole marker $[\cdot]^\tau$. For technical reasons it is important that this position is not syntactically restricted to just values: e.g., in a cell $x \text{ c } (k(\text{true}, \lambda y. e))$ the position of a hole can only be within e . A *process context* D is a process where exactly one process position is replaced with the hole marker $[\cdot]$.

We only consider well-typed contexts, where the typing of contexts is like the typing of expressions, with two additional typing rules for the context hole:

$$\frac{}{[\cdot]^\tau : \tau} \quad \frac{}{[\cdot] : wt}$$

The result of placing the expression $e : \tau$ (process p , resp.) in context C with hole $[\cdot]^\tau$ (context D , resp.), possibly capturing free variables of e (p , resp.), is written $C[e]$ ($D[p]$, resp.). It is easy to verify that $D[p] : wt$ if $D : wt$ and $p : wt$, and that for expression contexts $C[[\cdot]^\tau] : wt$ and expressions $e : \tau$, $C[e] : wt$ holds.

Fig. 5 defines *evaluation contexts* (ECs) E and *future ECs* F as particular contexts. ECs encode the standard call-by-value, left-to-right reduction strategy, while future ECs control dereferencing operations on futures and the triggering of suspended threads. The small-step reduction relation $p \rightarrow p'$ is the least binary relation on processes satisfying the rules in Fig. 6. We write $\xrightarrow{ev}$ for $\rightarrow$ when we want to distinguish reductions from *transformations* below. We use $\xrightarrow{*}$ ($\xrightarrow{+}$, resp.) for the reflexive-transitive (transitive, resp.) closure of $\rightarrow$. We sometimes label reductions with their name, e.g. $\xrightarrow{\text{BETA}(ev)}$, and $p \xrightarrow{a \vee b} q$ means that either $p \xrightarrow{a} q$ or $p \xrightarrow{b} q$.

Rule (CELL.NEW(ev)) creates new cells $z \text{ c } v$ with contents v . The exchange operation **exch**(z, v_1) writes v_1 to the cell and returns its previous contents. Since this is an atomic operation, no other thread can interfere. The rule (THREAD.NEW(ev)) spawns a new eager thread $x \leftarrow e$ where x may occur in e , so it may be viewed as a recursive declaration $x = e$. Similarly, (LAZY.NEW(ev)) creates a new suspended computation $x \xleftarrow{susp} e$. Dereferencing of future values (FUT.DEREF(ev)) and triggering of suspended computa-

Reduction rules.

(β -CBV(ev))	$E[(\lambda x. e) v] \rightarrow E[e[v/x]]$
(CELL.NEW(ev))	$E[\text{ref } v] \rightarrow (\nu z)(E[z] \mid z \text{ c } v)$
(CELL.EXCH(ev))	$E[\text{exch}(z, v_1)] \mid z \text{ c } v_2 \rightarrow E[v_2] \mid z \text{ c } v_1$
(THREAD.NEW(ev))	$E[\text{thread } v] \rightarrow (\nu z)(E[z] \mid z \leftarrow v z)$
(FUT.DEREF(ev))	$F[x] \mid x \leftarrow v \rightarrow F[v] \mid x \leftarrow v$
(LAZY.NEW(ev))	$E[\text{lazy } v] \rightarrow (\nu z)(E[z] \mid z \xleftarrow{susp} v z)$
(LAZY.TRIGGER(ev))	$F[x] \mid x \xleftarrow{susp} e \rightarrow F[x] \mid x \leftarrow e$
(HANDLE.NEW(ev))	$E[\text{handle } v] \rightarrow (\nu z)(\nu z')(E[v z z'] \mid z' \text{ h } z)$
(HANDLE.BIND(ev))	$E[x v] \mid x \text{ h } y \rightarrow E[\text{unit}] \mid y \leftarrow v \mid x \text{ h } \bullet$
(CASE.BETA(ev))	$E[\text{case } k_j(v_1, \dots, v_{ar(k_j)}) \text{ of } (k_i(x_1, \dots, x_{ar(k_i)}) \Rightarrow e_i)^{i=1 \dots n}]$ $\rightarrow E[e_j[v_1/x_1, \dots, v_{ar(k_j)}/x_{ar(k_j)}]]$

Well-formed processes. The rules can only be applied to well-formed processes.

Distinct variable convention. We assume that all processes to which rules apply satisfy the distinct variable convention, and that all new binders use fresh variables (z above). Reduction results will satisfy the distinct variable convention, if after β -CBV(ev), CASE.BETA(ev) and FUT.DEREF(ev) where values with bound variables may be copied, α -renaming is performed before applying the next rule.

Closure. Rule application is closed under structural congruence and process ECs D : if $p_1 \equiv D[p'_1]$, $p'_1 \rightarrow p'_2$, and $D[p'_2] \equiv p_2$ then $p_1 \rightarrow p_2$.

Figure 6. Small-step reduction relation $\rightarrow$ ($\xrightarrow{ev}$, resp.) of λ^τ (fch)

tions (LAZY.TRIGGER(ev)) is controlled by future evaluation contexts F . The rule (HANDLE.NEW(ev)) creates handle components. The application $x v$ in (HANDLE.BIND(ev)) “consumes” the handle x and binds y to v , resulting in a used handle $x \text{ h } \bullet$ and thread $x \leftarrow v$. In particular, notice that a used handle component $y \text{ h } \bullet$ indicates that y is a handle that has already been used to bind its associated future. It gets stuck if an expression $E[y v]$ has to be evaluated. Type safety holds for reduction, i.e., reduction preserves well-typedness (and well-formedness) of processes.

As an example, let $r: \text{ref List}(\tau)$ be a reference cell containing τ -lists. The effect of reducing process $r \text{ c } v \mid (\nu x)(x \leftarrow \text{thread } \lambda y. (\text{exch}(r, \text{cons}(v', y))))$ is the substitution of the cell content v by $\text{cons}(v', v)$ in a *single atomic step*. More precisely, it is $(\nu x)(\nu y)(r \text{ c } \text{cons}(v', y) \mid x \leftarrow y \mid y \leftarrow v)$. This process is observationally equivalent to $r \text{ c } \text{cons}(v', v)$.

As a second example, suppose **ack** is a (nullary) type constructor with a single data constructor **Ack**, and assume p is a thread suspending on a handled future x of type **ack**: $p \equiv F[\text{case } x \text{ of Ack} \Rightarrow e]$. The thread may resume the computation of e once a second thread uses the associated handler and provides a value for x : the process $p \mid F'[h \text{ Ack}] \mid h \text{ h } \bullet$ can reduce to $F[e] \mid F'[\text{unit}] \mid x \leftarrow \text{Ack} \mid h \text{ h } \bullet$, after (HANDLE.BIND(ev)), (FUT.DEREF(ev)) and (CASE.BETA(ev)) reductions. In this way, handled futures are the basic synchronization construct in λ^τ (fch).

Observations and contextual equivalence. A process p is *successful* if it is well-formed and in every component $x \leftarrow e$ of p , the identifier x is bound possibly via a chain $x \leftarrow x_1 \mid x_1 \leftarrow x_2 \mid \dots \mid x_{n-1} \leftarrow x_n \mid x_n \leftarrow v$ to a non-variable value, a cell or a lazy future, a handle, or a handled future. Hence, in a non-failing computation, every non-lazy future eventually refers to a “proper” value. For instance, $x \leftarrow \lambda y. y$,

$x \Leftarrow y \mid y \Leftarrow \langle x, x \rangle$ and $x \Leftarrow y \mid y \Leftarrow z$ are successful, while $x \Leftarrow x$ (a black hole) and $x \Leftarrow (\lambda u. \lambda v. v) (y \text{ unit}) \mid y \Leftarrow (\lambda u. \lambda v. v) (x \text{ unit})$ (a deadlocked process) are ruled out.

Reductions keep the type of expressions. Also, the following progress lemma holds: A well-formed and well-typed closed process either can be reduced, or it is successful, or it is not successful since it has a cycle of futures, or it is not successful, since it is deadlocked, where we use the following definitions: A *waiting thread* is of the form $x \Leftarrow \tilde{F}[y]$ and no reduction can be applied; a *finished thread* is of the form $x \Leftarrow v$; and a *deadlocked process* is a process where every thread component is finished or waiting, and there is at least one waiting thread.

We use $p \Downarrow$ to indicate that p is *may-convergent*, i.e., that there is a sequence of reductions $p \rightarrow^* p'$ such that p' is successful, and $p \Downarrow$ if the process is *must-convergent*, meaning that all reduction descendants p' of p are may-convergent. Dually, we call p *must-divergent* ($p \Uparrow$) if it has no reduction descendant that succeeds, and *may-divergent* ($p \uparrow$) if some reduction descendant of p is must-divergent. Thus, $p \Uparrow \Leftrightarrow \neg p \Downarrow$ and $p \uparrow \Leftrightarrow \neg p \Downarrow$. For $\xi \in \{\downarrow, \Downarrow\}$, we define contextual approximations between processes p_1, p_2 by:

$$p_1 \leq_\xi p_2 \Leftrightarrow \forall D : D[p_1] \xi \Rightarrow D[p_2] \xi$$

We write $p_1 \leq p_2$ if both $p_1 \leq_\downarrow p_2$ and $p_1 \leq_\Downarrow p_2$ hold, and $p_1 \sim p_2$ if both $p_1 \leq p_2$ and $p_2 \leq p_1$ hold. The same definitions for expressions e_1, e_2 of equal type τ and expression contexts $C[\cdot]$ results in relations $\leq_{\downarrow, \tau}$, $\leq_{\Downarrow, \tau}$, and $\sim_\tau$.

Translations. We recall the framework of Schmidt-Schauß et al. (2008), where an abstract calculus $\mathcal{C}$ consists of sets of (well-typed) processes p , contexts D , and convergence predicates ξ . The calculus $\lambda^\tau(\text{fch})$ and the other (possibly untyped) calculi introduced in the subsequent sections fit into this general framework. A *translation* T between two such calculi maps well-typed processes to well-typed processes, and contexts to contexts. A translation T between calculi $\mathcal{C}$ and $\mathcal{C}'$ is *convergence equivalent* if $T(p) \xi \Leftrightarrow p \xi$ for all p and all convergence predicates ξ . The translation T is *compositional* iff for all contexts D and processes p we have $T(D)[T(p)] = T(D[p])$. A translation T is *observationally correct*, if for all programs p and all contexts D , and all convergence predicates ξ : $T(D[p]) \xi \Leftrightarrow T(D)[T(p)] \xi$. A translation is *adequate* if T reflects operational approximation, i.e. if $T(p_1) \leq_{\mathcal{C}'} T(p_2) \Rightarrow p_1 \leq_{\mathcal{C}} p_2$ for all p_1, p_2 . Finally, if T additionally preserves inequations, i.e. if $T(p_1) \leq_{\mathcal{C}'} T(p_2) \Leftrightarrow p_1 \leq_{\mathcal{C}} p_2$ for all p_1, p_2 holds, then it is *fully abstract*.

As described in the introduction, adequacy and full abstraction relate to the equational theories of the source and target language of a translation, and adequate translations provide useful tools for transferring equations. The soundness of an encoding, in the sense that each program is indistinguishable from its translation, is given by observational correctness. Of course, these notions are related:

Proposition 2.1 (Adequacy, (Schmidt-Schauß et al. 2008)). *If a translation T is compositional and convergence equivalent, then T is adequate and observationally correct. Moreover, observational correctness of a translation implies adequacy of the translation.*

For the considered calculi in this paper we also require compositionality on expressions, i.e. $T(C[e]) = T(C)[T(e)]$ and type correctness of T for expressions. These requirements simplify the corresponding proofs for processes. Moreover, they enable us to derive equivalences from the adequacy of the translations not only between processes but also between expressions, i.e. that $T(e) \sim_{T(\tau)} T(e')$ implies $e \sim_\tau e'$.

It is easy to verify that translations compose:

Proposition 2.2 (Composition). *Let $\mathcal{C}, \mathcal{C}', \mathcal{C}''$ be 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*

$$\begin{aligned} (\text{FUT.DEREF(a)}) \quad & C[x] \mid x \Leftarrow v \rightarrow C[v] \mid x \Leftarrow v \\ (\beta\text{-CBV(a)}) \quad & C[(\lambda x. e) v] \rightarrow C[e[v/x]] \\ (\text{CASE.BETA(a)}) \quad & C[\text{case } k_j(v_1, \dots, v_{ar(k_j)}) \text{ of } (k_i(x_1, \dots, x_{ar(k_i)}) \Rightarrow e_i)^{i=1 \dots n}] \\ & \rightarrow C[e_j[v_1/x_1, \dots, v_{ar(k_j)}/x_{ar(k_j)}]] \\ (\text{CELL.DEREF}) \quad & p \mid y \Leftarrow x \mid x \Leftarrow v \rightarrow p \mid y \Leftarrow v \mid x \Leftarrow v \\ (\text{GC}) \quad & p \mid (\nu y_1) \dots (\nu y_n) p' \rightarrow p \\ & \text{if } p' \text{ is successful and } y_1, \dots, y_n \text{ contain all} \\ & \text{process variables of } p' \\ (\text{DET.EXCH}) \quad & (\nu x)(y \Leftarrow \tilde{E}[\text{exch}(x, v_1)] \mid x \Leftarrow v_2) \\ & \rightarrow (\nu x)(y \Leftarrow \tilde{E}[v_2] \mid x \Leftarrow v_1) \end{aligned}$$

No capturing. We assume that no variables are moved out of their scope or into the scope of some other binder, i.e., $fv(v) \cap bv(C) = \emptyset$.

Closure. Transformations are closed under $\equiv$ and D -contexts.

Figure 7. Correct transformation rules for $\lambda^\tau(\text{fch})$

also a translation, and if T, T' are compositional (observationally correct, adequate, fully-abstract, respectively), then also the composition $T' \circ T$ is compositional (observationally correct, adequate, fully-abstract, respectively).

We also recall a criterion for fully abstract translations, which can be used if only new primitives are added to a calculus $\mathcal{C}'$. The statement of this criterion in (Schmidt-Schauß et al. 2008) contains a flaw; the following corrected version is proved in the technical report (Schmidt-Schauß et al. 2009).

Proposition 2.3 (Full abstraction for extensions). *Let $\mathcal{C}, \mathcal{C}'$ be two calculi, let $\iota : \mathcal{C}' \rightarrow \mathcal{C}$ (the embedding) and $T : \mathcal{C} \rightarrow \mathcal{C}'$ be compositional and convergence equivalent translations, such that $T \circ \iota$ is the identity on $\mathcal{C}'$ -programs, on $\mathcal{C}'$ -observers, and on $\mathcal{C}'$ -types. Then ι is fully abstract. If T is injective on types, then T is also fully abstract.*

If a translation T is observationally correct and injective on types, then $\leq$ is retained under T , relative to its image.

Remark 2.4 (on full abstraction on images). *A variation of this full abstraction result is possible (Schmidt-Schauß et al. 2009). Let $\mathcal{C}, \mathcal{C}'$ be calculi and $T : \mathcal{C} \rightarrow \mathcal{C}'$ be an observationally correct translation. Let $\mathcal{C}'' := T(\mathcal{C})$ be the subcalculus of $\mathcal{C}'$ consisting of the images under T , and let $\leq_T$ be the preorder defined on $\mathcal{C}''$. Moreover, assume that for all τ , T is surjective on the programs of type τ and for every τ' , T is a surjective mapping $T : \mathcal{O}_{\tau_1, \tau_2} \rightarrow \mathcal{O}_{T(\tau_1), T(\tau_2)}$, where $\mathcal{O}_{\tau_1, \tau_2}$ are the contexts of type τ_2 with hole of type τ_1 . Then for all types τ and programs p_1, p_2 : $p_1 \leq_\tau p_2 \Leftrightarrow T(p_1) \leq_{T(\tau)} T(p_2)$. That is, the translation is fully abstract as translation $T : \mathcal{C} \rightarrow \mathcal{C}''$.*

3. Correctness of Transformations in $\lambda^\tau(\text{fch})$

We will make use of program transformations, which are called *correct* if whenever p_1 is transformed into p_2 , then $p_1 \sim p_2$. In Fig. 7 some program transformations are defined. It is important that program transformations preserve the types of replaced subexpressions. E.g. the rule $(\beta\text{-CBV(a)})$ may also be applied from right to left, and in this case, we must choose a variable x with $\Gamma(x) = \tau$ where τ is the (uniquely determined) type of the value v . The use of the framework sketched at the end of Section 2 makes it possible, via translations, to lift process equivalences from the untyped lambda calculus with futures (Niehren et al. 2007) to correct pro-

gram transformations in $\lambda^\tau(\text{fch})$. We establish these correctness results by using an adequate translation from $\lambda^\tau(\text{fch})$ into $\lambda(\text{fh})$.

We first restate known equivalences for $\lambda(\text{fh})$. The calculus $\lambda(\text{fh})$ is the subcalculus of $\lambda^\tau(\text{fch})$ without constructors, case-expressions, and typing (see Niehren et al. 2007, for details).

Theorem 3.1. *All reductions of $\lambda(\text{fh})$ except $\text{CELL.EXCH}(\text{ev})$, and the transformations $\beta\text{-CBV}(\text{a})$, $\text{FUT.DEREF}(\text{a})$, CELL.DEREF , GC and DET.EXCH (see Fig. 7) are correct for $\lambda(\text{fh})$.*

Proof sketch. Correctness of the transformations was established in (Niehren et al. 2007). Although the standard reduction differs slightly from the one used here, we have show that this makes no difference for contextual equivalence in the technical report Schmidt-Schauß et al. (2008), and thus the statement holds. $\square$

Removing constructors and types. In order to lift contextual equivalences from $\lambda(\text{fh})$ to $\lambda^\tau(\text{fch})$, we construct a translation $T_C : \lambda^\tau(\text{fch}) \rightarrow \lambda(\text{fh})$ that is adequate. For the adequacy of this encoding it is necessary that $\lambda^\tau(\text{fch})$ is typed; otherwise untyped programs that get stuck due to a dynamic type error could become must-convergent after translation. Conversely, it is not possible to restrict $\lambda(\text{fh})$ to simple types, since the encoding of case and constructors cannot be monomorphically typed. These problems are discussed by Schmidt-Schauß et al. (2008), where illustrating examples can be found.

The main part in giving the translation is to encode **case** and constructors, and to show that the translation has all required properties. The encoding is a variant of the classic Church encodings. Let $K = \mathcal{D}(\kappa)$ be the set of constructors for a specific type constructor κ and let $\tau_1, \dots, \tau_m$ be types. By the assumptions on the signature, K is non-empty. For the definition of T_C we choose an arbitrary (but from now on fixed) order of the constructors in K , $k_1, \dots, k_n$ where $n \geq 1$. The two key cases of the encoding T_C are the following, where $l = \text{ar}(k_i)$:

$$\begin{aligned} T_C(k_i(e_1, \dots, e_l)) &\triangleq \\ \text{let } x_1 = T_C(e_1), \dots, x_l = T_C(e_l) \text{ in } \lambda p_1, \dots, p_n. p_i x_1 \dots x_l \text{ unit} \\ T_C(\text{case}_{\kappa} e \text{ of } (k_i(x_{i,1}, \dots, x_{i,\text{ar}(k_i)}) \Rightarrow e_i^{i=1 \dots n})) &\triangleq \\ T_C(e) (\lambda x_{1,1}, \dots, x_{1,\text{ar}(k_1)}. \lambda \dots. T_C(e_1)) & \\ \dots (\lambda x_{n,1}, \dots, x_{n,\text{ar}(k_n)}. \lambda \dots. T_C(e_n)) & \end{aligned}$$

and it is extended homomorphically to all other cases (only the translation of reference cells requires some additional care, to ensure that the translation of the stored value is again syntactically in the form of a value). The additional **unit** argument to p_i in the encoding of constructors achieves the correct behavior in the case of nullary constructors, with respect to call-by-value semantics. Correspondingly, in the encoding of **case**, the additional final abstraction in each branch leads to a uniform translation also for nullary constructors. Finally, note that types are removed by this encoding, i.e. $\text{thread}^\tau \triangleq \text{thread}$, $\text{ref}^\tau \triangleq \text{ref}$, $\text{lazy}^\tau \triangleq \text{lazy}$, and $\text{handle}^\tau \triangleq \text{handle}$.

Proposition 3.2 (Adequacy). *The translation T_C is adequate.*

Remark 3.3. *Note that T_C is not fully abstract. We give an example without proof: The expressions $\lambda x^{\text{bool}}.x$ and $(\lambda x^{\text{bool}}. \text{if } x \text{ then } x \text{ else } x)$ are equivalent, but the translated expressions are not equivalent since they behave differently when applied to **unit**. The second expression is translated into $\lambda x.x$ ($\lambda \dots.x$) ($\lambda \dots.x$). Applying both expressions to **unit** will result in **unit** and in $(\text{unit } (\lambda \dots.\text{unit}) (\lambda \dots.\text{unit}))$, respectively. The first is a value, and the second is must-divergent. In particular, this shows that Proposition 2.3 is not applicable: there is no identity embedding from $\lambda(\text{fh})$ into $\lambda^\tau(\text{fch})$, since the former is untyped.*

Syntactic extensions:

$$\begin{aligned} \tau &\in \text{Type} ::= \text{buf } \tau \mid \dots \\ c &\in \text{Const} ::= \text{buffer}^\tau \mid \text{get}^\tau \mid \dots \\ e &\in \text{Exp} ::= \text{put}(e_1, e_2) \mid \dots \\ p &\in \text{Proc} ::= x \text{ b } - \mid x \text{ b } v \mid \dots \\ \tilde{E} &::= \text{put}(\tilde{E}, e) \mid \text{put}(v, \tilde{E}) \mid \dots \\ \tilde{F} &::= \tilde{E}[\text{put}([], v)] \mid \tilde{E}[\text{get} []] \mid \dots \end{aligned}$$

Extensions of the type system:

$$\frac{\tau \preceq \text{unit} \rightarrow \text{buf } \alpha}{\text{buffer}^\tau : \tau} \quad \frac{\tau \preceq \text{buf } \alpha \rightarrow \alpha}{\text{get}^\tau : \tau} \quad \frac{e_1 : \text{buf } \tau \quad e_2 : \tau}{\text{put}(e_1, e_2) : \text{unit}}$$

$$\frac{}{x \text{ b } - : wt} \quad \frac{}{x \text{ b } v : wt}$$

Extensions of the reduction rules:

$$\begin{aligned} (\text{BUFF.NEW}(\text{ev})) \ E[\text{buffer } v] &\rightarrow (\nu x)(E[x] \mid x \text{ b } -) \text{ fresh } x \\ (\text{BUFF.PUT}(\text{ev})) \ E[\text{put}(x, v)] \mid x \text{ b } - &\rightarrow E[\text{unit}] \mid x \text{ b } v \\ (\text{BUFF.GET}(\text{ev})) \ E[\text{get } x] \mid x \text{ b } v &\rightarrow E[v] \mid x \text{ b } - \end{aligned}$$

Figure 8. Extensions of $\lambda^\tau(\text{fch})$ for $\lambda^\tau(\text{fchb})$

Correctness of program transformations for $\lambda^\tau(\text{fch})$ follows by encoding the transformations into $\lambda(\text{fh})$ using T_C and then applying adequacy of T_C . This gives us the main result of this section:

Theorem 3.4. *The reduction rules of $\lambda^\tau(\text{fch})$, except for $\text{CELL.EXCH}(\text{ev})$, and the transformations of Fig. 7 (note the arbitrary contexts C in the first three) are correct for $\lambda^\tau(\text{fch})$.*

4. Concurrent Buffers are Encodable in $\lambda^\tau(\text{fch})$

By extending the syntax and operational semantics of $\lambda^\tau(\text{fch})$, we provide a specification of one-place buffers that describes their desired behavior.

The calculus $\lambda^\tau(\text{fch})$ is extended by new primitives for concurrent buffers. This defines the calculus $\lambda^\tau(\text{fchb})$, with the syntactic extensions shown in Fig. 8. $\lambda^\tau(\text{fchb})$ has two new components: $x \text{ b } -$ which represents an empty buffer, and $x \text{ b } v$ which represents a buffer that contains the value v . There are new constants **buffer** $^\tau$ to create a new buffer and **get** $^\tau$ to obtain the contents of a non-empty buffer (and emptying the buffer). There is also a new binary operator **put**, to place a new value into an empty buffer. Contexts C are as before, but extended to the new syntax, such that exactly one expression-position, which is not restricted to values, is replaced with a typed hole marker $[\cdot]^\tau$.

Fig. 8 also summarizes the operational interpretation of the new constructs, and extends the set of (future) evaluation contexts. Note that the reduction rules entail that **get** x suspends on an empty buffer x while **put** (x, v) suspends on a non-empty x . For typing we assume a new type constructor **buf** of arity 1. The typing of the constants is given by (instances of) type schemes (see Fig. 8); type safety then extends to the calculus $\lambda^\tau(\text{fchb})$. Contextual preorder is defined as expected: the notion of a *successful process* from $\lambda^\tau(\text{fch})$ is extended so that $\lambda^\tau(\text{fchb})$ also allows $x \text{ b } -$ and $x \text{ b } v$ as components of successful processes. A $\lambda^\tau(\text{fchb})$ process is *well-formed* if (in addition to the other process variables) no buffer variables are introduced twice.

In the remainder of this section we will show that there is an observationally correct translation $T_B : \lambda^\tau(\text{fchb}) \rightarrow \lambda^\tau(\text{fch})$ which implements buffers by handled futures. The proof of observational correctness of T_B requires equivalences in $\lambda^\tau(\text{fch})$, which have been derived in Theorem 3.4.

```

buffer  $\triangleq$   $\lambda_{-}$ . let  $\langle h, f \rangle = \text{newhandled}, \langle h', f' \rangle = \text{newhandled},$ 
      putg = ref(true), getg = ref(f),
      stored = ref(f'), handler = ref(h)
(1)    in thread  $\lambda_{-}.\langle \text{putg}, \text{getg}, \text{stored}, \text{handler} \rangle$  end

put  $\triangleq$   $\lambda\langle x_p, x_g, x_s, x_h \rangle, v$ . let  $\langle h, f \rangle = \text{newhandled}$ 
(1)    in wait (exch( $x_p, f$ ));
(2)    exch( $x_s, v$ );
(3)    (exch( $x_h, h$ ))(true) end

get  $\triangleq$   $\lambda\langle x_p, x_g, x_s, x_h \rangle.$ 
      let  $\langle h, f \rangle = \text{newhandled}; \langle h', f' \rangle = \text{newhandled}$ 
(1)    in wait (exch( $x_g, f$ ));
(2)    let  $v = (\text{exch}(x_s, f'))$ 
(3)    in (exch( $x_h, h$ ))(true);  $v$  end
end

```

Figure 9. Implementing the buffer operations *buffer*, *put* and *get*, where (1), (2), (3) indicate subexpressions for later reference.

4.1 Implementing Buffers Using Handled Futures

Any concrete realization of buffers will rely on (more or less intricate) non-interference properties and the preservation of various invariants. We consider a particular implementation of buffers in $\lambda^\tau(\text{fch})$, in terms of reference cells and handled futures. This induces a translation from $\lambda^\tau(\text{fchb})$ into $\lambda^\tau(\text{fch})$.

The implementation in $\lambda^\tau(\text{fch})$ of operations corresponding to **buffer**, **put**, and **get** is shown in Fig. 9. The buffer data structure is implemented as a tuple, consisting of four reference cells:

$\text{buf } \tau \triangleq \text{ref bool} \times \text{ref bool} \times \text{ref } \tau \times \text{ref (bool} \rightarrow \text{unit)}.$

The first and second of these reference cells serve as guards to ensure that succeeding *put* and *get* operations alternate. Exactly one of them will contain a handled future: if the first guard contains a future, this indicates that the buffer is currently non-empty, hence *put* must block. Likewise, if the second guard contains a handled future, the tuple represents an empty buffer and *get* must block. The final reference cell stores a handler for this future. The third cell, of type $\text{ref } \tau$, stores the actual contents of the buffer. When representing an empty buffer, this reference will contain a handled future of type τ as a dummy value. In summary, there are the following invariants associated with the value $\langle \text{putg}, \text{getg}, \text{stored}, \text{handler} \rangle$:

- the guards *putg* and *getg* contain either a handled future or true (perhaps reachable via dereferencing futures),
- at most one of *putg* and *getg* contains true,
- if *getg* contains true then the value in *stored* is the value currently in the buffer,
- whenever *putg* contains true then the value in *stored* is ‘garbage’, representing an empty buffer.

The procedure *buffer* yields a tuple representing an empty buffer, satisfying the invariants. The procedure *put*, when applied to a buffer $\langle \text{putg}, \text{getg}, \text{stored}, \text{handler} \rangle$ and a value v , suspends until the buffer is guaranteed to be empty. This is achieved by pattern matching on the contents of *putg* (using *wait*): since the first argument position of the **case** construct constitutes a future EC, *put* can continue only when *putg* contains a proper (non-future) value. By the invariants, this implies that the buffer is empty. At the same time, *putg* is replaced by a fresh future f , with handle h , to indicate that the buffer will be non-empty after *put* succeeds. After writing v to the cell *stored*, the second guard *getg* is set to true (perhaps

via a reference) to permit following *get* operations to succeed. This is done using the handle stored in the reference cell *handler*, which is replaced by the handle h for the freshly introduced future f . The procedure *get* is analogous (partly symmetric) to *put*.

The use of the handled futures in *put* and *get* is somewhat subtle: in general, several threads concurrently attempt to place values into the buffer (and dually, for reading from the buffer). The thread that is scheduled first replaces the contents of the guard by a future f_1 . This future can be bound only after this instance of *put* has finished. A second instance of *put* can proceed immediately with its own exchange operation, replacing f_1 by a future f_2 before the wait suspends on f_1 . In this way, a chain of threads suspending on futures $f_1, f_2, \dots$ in their respective *put* operations can build up. At the same time, a chain of threads suspending in their respective *get* operations can build up.

4.2 Implementation as Translation

The implementation gives rise to a translation T_B from $\lambda^\tau(\text{fchb})$ into $\lambda^\tau(\text{fch})$: **put**, **get**, and **buffer** are replaced by the resp. program code, *put*, *get*, and *buffer* from Fig. 9, where for **put**, the two arguments are translated into a pair. On process level, we replace:

$$\begin{aligned}
T_B(x \text{ b } -) &\triangleq (\nu x_p)(\nu x_g)(\nu x_s)(\nu x_h) x \Leftarrow \langle x_p, x_g, x_s, x_h \rangle \\
&\quad | (\nu h)(\nu f)(\nu h')(\nu f')(h \text{ h } f \mid h' \text{ h } f' \mid x_p \text{ c true} \\
&\quad \mid x_g \text{ c } f \mid x_s \text{ c } f' \mid x_h \text{ c } h) \\
T_B(x \text{ b } v) &\triangleq (\nu x_p)(\nu x_g)(\nu x_s)(\nu x_h) x \Leftarrow \langle x_p, x_g, x_s, x_h \rangle \\
&\quad | (\nu h)(\nu f)(h \text{ h } f \mid x_p \text{ c } f \mid x_g \text{ c true} \\
&\quad \mid x_s \text{ c } T_B(v) \mid x_h \text{ c } h).
\end{aligned}$$

Formally, these replacements homomorphically extend to a mapping $T_B : \lambda^\tau(\text{fchb}) \rightarrow \lambda^\tau(\text{fch})$ on all $\lambda^\tau(\text{fchb})$ -expressions, -processes, and -contexts. A corresponding type translation is defined inductively, by $T_B(\text{buf } \tau) \triangleq \text{buf}(T_B(\tau))$, and proceeding homomorphically in all other cases. Note that T_B is not injective on the types, since buffer-types are mapped to the type of 4-tuples. These mappings are compatible with typing:

1. If $e : \tau$ then $T_B(e) : T_B(\tau)$.
2. If p is well-typed, then $T_B(p)$ is well-typed.
3. If p is well-formed, then $T_B(p)$ is well-formed.
4. For $C[[\cdot]^\tau] : wt$ and $D : wt$ we have $T_B(C[[\cdot]^{T_B(\tau)}]) : wt$ and $T_B(D) : wt$.

Corresponding typing properties also hold for contexts, so that T_B forms a translation in the sense of Section 2 that is compositional:

Lemma 4.1 (Compositionality). *The translation $T_B : \lambda^\tau(\text{fchb}) \rightarrow \lambda^\tau(\text{fch})$ is compositional, i.e., for all p, D we have $T_B(D)[T_B(p)] = T_B(D[p])$, and for all τ , and $e : \tau$ and $C[[\cdot]^\tau]$, we have $T_B(C)[T_B(e)] = T_B(C[e])$.*

Proof. Immediate from the fact that T_B is extended homomorphically from constants to all terms, and from base components to arbitrary processes, resp. $\square$

We argue that the buffer implementation, described by T_B above, is correct. To this end, we will prove T_B convergence equivalent in this section, and use compositionality (Lemma 4.1). By Proposition 2.1, this entails the observational correctness of T_B .

Lemma 4.2 (T_B preserves success). *Let p be a $\lambda^\tau(\text{fchb})$ -process.*

1. If p is successful, then so is $T_B(p)$. In particular, $T_B(p) \Downarrow$ in this case.
2. If $T_B(p)$ is successful, then p is also a successful process.

The definition of the translation T_B also shows the following.

Lemma 4.3. *If D is a process context of $\lambda^\tau(\text{fchb})$, then $T_B(D)$ is a process context. If E is an evaluation context of $\lambda^\tau(\text{fchb})$, then $T_B(E)$ is an evaluation context in $\lambda^\tau(\text{fch})$.*

Note that a corresponding property does not hold for future evaluation contexts. As an example, consider the process $x \leftarrow \text{put}(y, v) \mid y \leftarrow \text{get } x \mid \dots$. Assume that **get** is executed first, then **put**. For the corresponding reduction sequence in $\lambda^\tau(\text{fch})$ it is unavoidable that essentially the same sequence is used on the implementation *get* and *put*. However, the initial reductions of *put* may be executed earlier. (In the case of $y \xrightarrow{\text{susp}} \text{get } x$, this is even enforced.) For the reduction in the implementation this means that the reduction steps of instances of *get* and *put* cannot be gathered into one contiguous block; this is possible only for the main steps 1,2,3 of an instance.

4.3 Observational Correctness of the Translation T_B

Proposition 4.4 ($\downarrow$ -preservation of T_B). *For every $\lambda^\tau(\text{fchb})$ -process p , $p \downarrow \Rightarrow T_B(p) \downarrow$.*

Proof sketch. The proof can be done by induction on the length of a reduction U of p , where one has to take into account that the future evaluation contexts are not preserved, and so the sequence of reductions is not straightforward. The idea is to construct a sequence of reductions and transformations in $\lambda^\tau(\text{fch})$, and then use Theorem 3.4 that shows that we have only inserted correct $\lambda^\tau(\text{fch})$ -transformations. $\square$

The other parts of the proof of convergence equivalence of T_B require a more careful analysis.

Invariants of the Buffer-Implementation. We sketch an analysis of the global state of the $\lambda^\tau(\text{fch})$ -translation $T_B(p)$ during reduction. In this analysis we will focus on the reductions of *put* and *get*. The reduction of *buffer* provides no problems, since all reductions are correct in $\lambda^\tau(\text{fch})$ according to Theorem 3.4.

The execution of each instance of *put* or *get* consists of initial (β -CBV(ev)) and (CASE.BETA(ev)) reductions, and eventually the argument has to be evaluated in a future-strict context.

It is possible to analyze exactly the global state and the changes that belong to one buffer x . We will only informally describe the results here. The analysis has to take into account all the active instances of *put* and *get*, i.e., the code-pointer and internal state of the implemented *puts* and *gets* until they are finished. Therefore, we require the code-pointer that may have values 1a, 1b, 2, 3a, 3b as in the encoding in Fig. 9, where the indexing is according to the sequential execution. For the **put**-encoding:

- (1) wait (**exch**(x_p, f)); (1a) for **exch** ...; (1b) for wait
- (2) **exch**(x_s, v_i);
- (3) (**exch**($x_h, h_{p,i}$))(true) (3a) for **exch** ...; (3b) for handle-binding

For the **get**-encoding:

- (1) wait (**exch** $x_g f$); (1a) for **exch** ...; (1b) for wait
- (2) let $v = \text{exch}(x_s, f'_{g,i})$;
- (3) in (**exch**($x_h, h_{g,i}$))(true); v (3a) for **exch** ...; (3b) for handle-binding

Informally, there are two queues: one queue of *put*-instances blocking on guards that need to be activated by the respective handlers by the previous *get*, and another queue of *get*-instances that wait for activation of their guards by handlers of the previous *put*. All the futures will be bound to true at some time, provided the evaluation terminates successfully.

Using induction on the number of buffer-related small-step reductions, we see that the above description is an invariant for the buffer-implementation of a single buffer x . Our analysis implies the

following sequencing constraint of reductions in $\lambda^\tau(\text{fch})$, where $\xrightarrow{x,a,b}$ means a reduction step for a particular buffer (implementation) x of the *put/get*-instance b with code-pointer a .

Lemma 4.5. *For a fixed buffer x the following sequence relations holds in any $\xrightarrow{\text{ev}}$ -reduction U :*

*If for two instances b_1, b_2 of *put, get*: $\xrightarrow{x,a,b_1}$ is before $\xrightarrow{x,a,b_2}$ for some $a \in \{1b, 2, 3a, 3b\}$, then for all $a_1, a_2 \in \{1b, 2, 3a, 3b\}$: $\xrightarrow{x,a_1,b_1}$ is before $\xrightarrow{x,a_2,b_2}$.*

*For two *put*-instances b_1, b_2 (or *get*-instances, respectively): $\xrightarrow{x,1a,b_1}$ is before $\xrightarrow{x,1a,b_2}$ iff $\xrightarrow{x,2,b_1}$ is before $\xrightarrow{x,2,b_2}$.*

Proposition 4.6 ($\downarrow$ -reflection of T_B). *For every $\lambda^\tau(\text{fchb})$ -process p , $T_B(p) \downarrow \Rightarrow p \downarrow$.*

Proof sketch. The main idea of the proof is to rearrange a reduction sequence U of $T_B(p)$ using Lemma 4.5 until the reduction steps that belong to an instance of *put/get* are in a contiguous block. Since future evaluation contexts do not correspond before and after the translation, we also have to use the equivalences from Theorem 3.4 for further rearrangement. $\square$

Proposition 4.7 ($\Downarrow$ -reflection of T_B). *For every $\lambda^\tau(\text{fchb})$ -process p , $T_B(p) \Downarrow \Rightarrow p \Downarrow$.*

Proof. Suppose that for the $\lambda^\tau(\text{fchb})$ -process p we have $p \uparrow$. We show $T_B(p) \uparrow$. Since $p \uparrow$ there is a reduction R from p to a process $p_0 \uparrow$. Analogous to the proof of Proposition 4.4, we can show by induction on the length of R that there is a sequence R' of correct transformations and reductions from $T_B(p)$ to the process $T_B(p_0)$. Proposition 4.6 applied to p_0 shows that $T_B(p_0) \downarrow$ is impossible, hence $T_B(p_0) \uparrow$ holds. By induction on the length of R' (which consists of ev-reductions and correct transformations), Theorem 3.4 is used to show $T_B(p) \uparrow$. $\square$

The proof of the following proposition is more intricate:

Proposition 4.8 ($\Downarrow$ -preservation of T_B). *For all $p \in \lambda^\tau(\text{fchb})$: $p \Downarrow \Rightarrow T_B(p) \Downarrow$.*

Proof. The detailed proof is in (Schwinghammer et al. 2009); here we give a sketch. We prove the equivalent claim that for every $\lambda^\tau(\text{fchb})$ -process p , $T_B(p) \uparrow \Rightarrow p \uparrow$. As in the proof of Proposition 4.6, a given reduction R corresponding to $T_B(p) \uparrow$ will be rearranged and modified in order to construct a $\lambda^\tau(\text{fchb})$ -reduction of p that shows $p \uparrow$. This allows a similar rearrangement into 1,2,3-blocks and intermediate correct transformations. However, this is not possible for all instances of *put* and *get*, since some of these may be started but never completed in the reduction.

Variants of the following argument can be used to overcome this difficulty. Suppose a certain instance m of *put* has been started within R , but the next reduction step, say from (3a), is missing in R . Let p_ω be the last process in R , for which we necessarily have $p_\omega \uparrow$. The commutation properties show that $\xrightarrow{3a,m}$ is a reduction possibility of p_ω , i.e. $p_\omega \xrightarrow{3a,m} p_{\omega,1}$, which immediately implies $p_{\omega,1} \uparrow$. Thus we extend R to $R \cdot \xrightarrow{3a,m}$. This procedure is repeated until all partially executed instances are completed; as an invariant, the number of started instances of *put, get* is not increased. Now one can construct a reduction sequence showing $p \uparrow$. Finally, there is some process p_ω with $p \xrightarrow{*} p_\omega$ and $T_B(p_\omega) = q_\omega$, where q_ω is the final process of the rearranged and extended sequence R' . The property $q_\omega \uparrow$ can be shown using Theorem 3.4. Lemma 4.4 shows that $p_\omega \downarrow$ is impossible, hence $p_\omega \uparrow$. Thus, $p \uparrow$. $\square$

Propositions 4.4, 4.6, 4.7, 4.8, and 2.1 imply:

Theorem 4.9 (Observational correctness of T_B). *The translation T_B is convergence equivalent. In fact, it is observationally correct: for all C and e , $C[e]$ and $T_B(C)[T_B(e)]$ have the same convergence behavior.*

As a consequence of this theorem, T_B is also adequate. For the proofs in the next section, we need also the correctness of several transformations in $\lambda^\tau(\text{fchb})$, which follow from adequacy.

Theorem 4.10 (Correct transformations in $\lambda^\tau(\text{fchb})$). *All reduction rules of $\lambda^\tau(\text{fchb})$ are correct, with the exception of $\text{CELL.EXCH}(\text{ev})$, $\text{BUFF.GET}(\text{ev})$, and $\text{BUFF.PUT}(\text{ev})$. The transformations $\beta\text{-CBV}(\mathbf{a})$, $\text{FUT.DEREF}(\mathbf{a})$, CELL.DEREF , GC and DET.EXCH (see Fig. 7) lifted to $\lambda^\tau(\text{fchb})$ are correct.*

Proof sketch. The statement follows by translating the transformations into the calculus $\lambda^\tau(\text{fch})$ using T_B . Then, the adequacy of T_B , some reasoning using Lemma 4.3, and Proposition 4.4 show the claim. $\square$

Although we do not make use of this fact, we can show that the embedding $\iota_B : \lambda^\tau(\text{fch}) \rightarrow \lambda^\tau(\text{fchb})$ is fully abstract, using Proposition 2.3.

Theorem 4.11 (Full abstraction of the embedding). *The embedding $\iota_B : \lambda^\tau(\text{fch}) \rightarrow \lambda^\tau(\text{fchb})$ is fully abstract.*

As an aside, note that we cannot use Proposition 2.3 to also show full abstraction of the translation T_B : the proposition requires injectivity on types, which T_B does not satisfy.

4.4 Applications of Observational Correctness

We can use the observational correctness of the translation T_B to derive a number of consequence for the implementation of buffers, and in turn obtain justification for considering this notion.

A common approach to the specification of abstract data types, in the sequential case, is by an axiomatic description of the operations. The results developed above allow us to prove that the buffers of $\lambda^\tau(\text{fchb})$ satisfy a number of such axioms. Using adequacy of T_B , the implied correctness of transformations for $\lambda^\tau(\text{fchb})$ (Theorem 4.10), and correctness of program transformations for $\lambda^\tau(\text{fch})$ (Theorem 3.4), one can show that the following transformations for **put** and **get** are correct:

$$\begin{aligned} (\text{DET.PUT}) \ (\nu x).E[\mathbf{put}(x, v)] \mid x \mathbf{b} - &\rightarrow (\nu x).E[\mathbf{unit}] \mid x \mathbf{b} v \\ (\text{DET.GET}) \ (\nu x).E[\mathbf{get} \ x] \mid x \mathbf{b} v &\rightarrow (\nu x).E[v] \mid x \mathbf{b} - \end{aligned}$$

These rules are like $\text{BUFF.PUT}(\text{ev})$ and $\text{BUFF.GET}(\text{ev})$, but restricted to sequentially used buffers. Then, $\mathbf{get}(\mathbf{put}(\text{buffer } u, v)) \sim v$ and similar equivalences follow.

More generally, we would like to show that the code in Fig. 9 correctly implements the specification of buffers, even if they are used in a non-sequential context. In other words, any use of buffers should give rise to the same observations, whether one computes with buffers abstractly using the specification, or concretely using the implementation. Informally, such a result states that the implementation as well as the specification can be considered as different realizations of an abstract data type of buffers. Since formally, the two live in different calculi, we use convergence equivalence (Theorem 4.9) of the translation T_B directly, rather than arguing by its adequacy as done for (DET.PUT) and (DET.GET) above.

Specifically, let $e : \tau'$ be any “client” making use of buffers: $e : \tau'$ is a $\lambda^\tau(\text{fchb})$ -program that may have free occurrences of the variables $b : \text{unit} \rightarrow \text{buf} \tau$, $p : \text{buf} \tau \times \tau \rightarrow \text{unit}$ and $g : \text{buf} \tau \rightarrow \tau$ but does not otherwise contain the buffer primitives, and τ and τ' are $\lambda^\tau(\text{fch})$ types. Such client programs are not affected by the encoding T_B induced by the implementation; we have $T_B(C) = C$ for the context C defined as $\text{let } \langle b, p, g \rangle = [] \text{ in } e$, so convergence

$$\begin{aligned} T_H(h \mathbf{h} \ f) &\triangleq (\nu f')(f \xrightarrow{\text{susp}} \text{let } v = \mathbf{get} \ f' \text{ in } \mathbf{put}(f', v); v) \\ &\quad \mid f' \mathbf{b} - \mid h \leftarrow \lambda z. \mathbf{put}(f', z)) \\ T_H(h \mathbf{h} \ \bullet) &\triangleq h \xrightarrow{\text{susp}} h \\ T_H(\mathbf{handle}) &\triangleq \lambda x. \text{let} \\ &\quad f' = \mathbf{buffer \ unit} \\ &\quad f = \mathbf{lazy} \ (\lambda_. \text{let } v = \mathbf{get} \ f' \text{ in } \mathbf{put}(f', v); v) \\ &\quad h = \mathbf{thread} \ (\lambda_. \lambda z. \mathbf{put}(f', z)) \\ &\quad \text{in } x \ f \ h \ \text{end} \\ T_H(p) &\triangleq \text{homomorphically wrt. the term structure} \end{aligned}$$

Figure 10. Encoding of handles using buffers

equivalence yields

$$C[\langle \mathbf{buffer}, \lambda \langle x, y \rangle. \mathbf{put}(x, y), \mathbf{get} \rangle] \xi \Leftrightarrow C[\langle \mathbf{buffer}, \mathbf{put}, \mathbf{get} \rangle] \xi$$

for all observations $\xi \in \{\downarrow, \Downarrow\}$.

5. Handled Futures are Encodable with Buffers

In this section we show that we can encode handled futures using buffers. Let $\lambda^\tau(\text{fcb})$ be the subcalculus of $\lambda^\tau(\text{fchb})$ where handles are removed. More precisely, in $\lambda^\tau(\text{fcb})$ the components $y \mathbf{h} \ x$, $\mathbf{h} \bullet$, and the constant **handle** are removed from the syntax of processes, expressions, evaluation contexts etc. Consequently, the reductions $\text{HANDLE.BIND}(\text{ev})$, and $\text{HANDLE.NEW}(\text{ev})$ are also dropped.

We show that there exists a translation $T_H : \lambda^\tau(\text{fchb}) \rightarrow \lambda^\tau(\text{fcb})$ which is observationally correct and fully abstract. The translation T_H is defined in Fig. 10. It does not change the types.

The idea of the translation is to simulate the synchronization effect of handles using the synchronization mechanism of one-place buffers. We consider the encoding of a handle component $y \mathbf{h} \ x$. An empty buffer represents the ability to bind the handled future x , i.e. binding of a handle consists in performing a put operation on the buffer. If the handled future x is accessed for the first time, then it becomes bound to the content of the filled buffer and another put-operation is performed on the buffer to ensure that the buffer remains full. The encoding of the handled future using lazy threads ensures that the possibility that a handle is not used in a successful reduction is translated into a successful reduction in the buffer implementation. Using (non-lazy) threads in the encoding would end up in a suspended get-operation, which is never successful.

The encoding of the constant **handle** generates the translated components of a handle when it is applied to an argument (modulo some $\text{BETA}(\text{ev})$ -reductions). **Note that equivalently we could have defined $h = \lambda z. \mathbf{put}(f', z)$ for the let-binding of h , since $\lambda z. \mathbf{put}(f', z) \sim_\tau \mathbf{thread} \ (\lambda_. \lambda z. \mathbf{put}(f', z))$ can be proved in $\lambda^\tau(\text{fcb})$, using the correct transformations of Lemma 5.3.**

The encoding of the used handle is different, compared to the result of the reduction of an encoded handler use. Nevertheless, since a $\text{HANDLE.BIND}(\text{ev})$ -operation on a used handle is not possible and leads to a must-divergent process, it is sufficient to introduce the process component $h \xrightarrow{\text{susp}} h$, which fails as soon as an encoded $\text{HANDLE.BIND}(\text{ev})$ -operation is performed. Moreover, after we have proved adequacy of T_H , we can verify that $h \xrightarrow{\text{susp}} h \sim h \mathbf{h} \ \bullet$ in $\lambda^\tau(\text{fchb})$ and $\lambda^\tau(\text{fch})$, respectively (see Remark 5.9). The $f \xrightarrow{\text{susp}} \mathbf{get} \ f'$ and the **lazy**-operator in the encoding are necessary, otherwise the future f would enforce the concurrent evaluation of $\mathbf{get} \ f'$ even if f does not occur in an E -context.

The translation T_H is compatible with typing, i.e. if p is well-typed and well-formed, then $T_H(p)$ is well-typed and well-formed.

Corresponding typing properties hold for contexts, so that T_H is a translation in the sense of Section 2.

Lemma 5.1 (Compositionality of T_H). *The translation $T_H : \lambda^\tau(\text{fchb}) \rightarrow \lambda^\tau(\text{fcb})$ is compositional, i.e., for all p, D, e, C , we have $T_H(D)[T_H(p)] = T_H(D[p])$ and $T_H(C)[T_H(e)] = T_H(C[e])$.*

We sketch the proof of observational correctness and adequacy. The following properties of T_H are easy to verify:

Proposition 5.2. *An expression $e \in \lambda^\tau(\text{fchb})$ is a $\lambda^\tau(\text{fchb})$ -value iff $T_H(e)$ is a $\lambda^\tau(\text{fcb})$ -value. A process $p \in \lambda^\tau(\text{fchb})$ is a successful $\lambda^\tau(\text{fchb})$ -process iff $T_H(p)$ is a successful $\lambda^\tau(\text{fcb})$ -process. For contexts of $\lambda^\tau(\text{fchb})$: $T_H(D)$ ($T_H(E)$, $T_H(F)$, resp.) is a D -context (E -context, F -context, resp.) for $\lambda^\tau(\text{fcb})$ iff D (E , F , resp.) is a D -context (E -context, F -context, resp.) for $\lambda^\tau(\text{fchb})$.*

Because of this proposition, the hard cases for proving convergence equivalence of T_H are the encodings of $\text{HANDLE.NEW}(\text{ev})$ - and $\text{HANDLE.BIND}(\text{ev})$ -reductions; all other reduction are inherited by the translation. We first lift program equivalences from $\lambda^\tau(\text{fchb})$ to $\lambda^\tau(\text{fcb})$. Let $\iota_H : \lambda^\tau(\text{fcb}) \rightarrow \lambda^\tau(\text{fchb})$ be the identity translation from $\lambda^\tau(\text{fcb})$ into $\lambda^\tau(\text{fchb})$. Obviously, ι_H is compositional and convergence equivalent, hence it is adequate. (In fact, we will prove that ι_H is fully-abstract at the end of this section.) As an immediate consequence of Theorem 4.10 and the adequacy of ι_H we obtain the following correct transformations.

Lemma 5.3 (Correct transformations in $\lambda^\tau(\text{fcb})$). *All reduction rules of $\lambda^\tau(\text{fcb})$ are correct, with the exception of $\text{CELL.EXCH}(\text{ev})$, $\text{BUFF.GET}(\text{ev})$, and $\text{BUFF.PUT}(\text{ev})$. The transformations $\beta\text{-CBV}(\mathbf{a})$, $\text{FUT.DEREF}(\mathbf{a})$, CELL.DEREF , GC and DET.EXCH (see Fig. 7) lifted to $\lambda^\tau(\text{fcb})$ are correct.*

Proposition 5.4 ($\downarrow$ -preservation of T_H). *For all $\lambda^\tau(\text{fchb})$ -processes p , the following holds: if $p \downarrow$, then $T_H(p) \downarrow$.*

Proof. The proof is by induction on the length of a successfully ending reduction for p . The induction base is covered by Proposition 5.2. For the induction step let $p \xrightarrow{\text{ev}} p'$. As induction hypothesis we use that $T_H(p') \downarrow$. Due to the properties of the context translation (Proposition 5.2) it is easy to see that all reductions of $\lambda^\tau(\text{fchb})$ except for $\text{HANDLE.BIND}(\text{ev})$ and $\text{HANDLE.NEW}(\text{ev})$ can be transferred to the encoding in $\lambda^\tau(\text{fcb})$, i.e. if $p \xrightarrow{a, \text{ev}} p'$ with $a \notin \{\text{HANDLE.BIND}(\text{ev}), \text{HANDLE.NEW}(\text{ev})\}$ then $T_H(p) \xrightarrow{a, \text{ev}} T_H(p')$. Hence, for these cases we have $T_H(p) \downarrow$.

For $p \xrightarrow{\text{HANDLE.BIND}(\text{ev})} p'$ we have:
 $T_H(p) \xrightarrow{\text{FUT.DEREF}(\text{ev})} \xrightarrow{\beta\text{-CBV}(\text{ev})} \xrightarrow{\text{BUFF.PUT}(\text{ev})} \xrightarrow{sp_1} T_H(p')$, where sp_1 is a program transformation in $\lambda^\tau(\text{fcb})$ defined as:

$$\begin{aligned} (\nu f')(y) &\xleftarrow{\text{susp}} \text{let } v = \text{get } f' \text{ in put}(f', v); v \text{ end} \\ &\quad | f' \text{ b } w \mid x \leftarrow \lambda z. \text{put}(f', z)) \\ &\rightarrow y \leftarrow w \mid x \xleftarrow{\text{susp}} x \end{aligned}$$

Since these transformations are either ev -reductions or correct (correctness of (sp_1) is proved in (Schwinghammer et al. 2009)), $T_H(p') \downarrow$ implies $T_H(p) \downarrow$.

For $p \xrightarrow{\text{HANDLE.NEW}(\text{ev})} p'$ it holds:
 $T_H(p) \xrightarrow{\beta\text{-CBV}(\text{ev})} \xrightarrow{\text{BUFF.NEW}(\text{ev})} \xrightarrow{\beta\text{-CBV}(\text{ev})} \xrightarrow{\text{LAZY.NEW}(\text{ev})} \xrightarrow{\beta\text{-CBV}(\text{ev})} \xrightarrow{\text{THREAD.NEW}(\text{ev})} \xrightarrow{\beta\text{-CBV}(\text{ev})} \xrightarrow{\beta\text{-CBV}(\text{ev})} \xrightarrow{\beta\text{-CBV}(f)} T_H(p')$

where $\beta\text{-CBV}(f)$ is the following restriction of $\beta\text{-CBV}(\mathbf{a})$: The expression context C in the transformation must be flat, i.e. the hole marker is not below a binding construct for expressions (λ -binder, or pattern in a **case**-expression). Since all these transformations are correct (Lemma 5.3), we have $T_H(p') \downarrow$ implies $T_H(p) \downarrow$. $\square$

Proposition 5.5 ($\downarrow$ -reflection of T_H). *For all $\lambda^\tau(\text{fch})$ -processes p , the following holds: if $T_H(p) \downarrow$, then $p \downarrow$.*

Proof. Let $p \in \lambda^\tau(\text{fch})$ and $T_H(p) \downarrow$. We show how to derive a successful reduction sequence for p . We track the encoded handles and handle operations in the image of T_H (e.g. by using an appropriate labeling) and use induction on the length of the given reduction sequence for $T_H(p)$. If $T_H(p)$ is successful, then Proposition 5.2 implies that p must be successful, too. For the induction step let $T_H(p) \xrightarrow{\text{ev}} q \xrightarrow{\text{ev}, n} q'$ where q' is successful. There are three cases:

- The reduction $T_H(p) \xrightarrow{\text{ev}} q$ is not the first reduction of an encoded $\text{HANDLE.NEW}(\text{ev})$ or $\text{HANDLE.BIND}(\text{ev})$ -reduction. For these cases it holds $p \xrightarrow{\text{ev}} p'$ and $T_H(p') = q$.
- The reduction $T_H(p) \xrightarrow{\text{ev}} q$ is the first reduction of an encoded $\text{HANDLE.BIND}(\text{ev})$ reduction, i.e. it is a $\text{FUT.DEREF}(\text{ev})$ -reduction and $p \xrightarrow{\text{HANDLE.BIND}(\text{ev})} p'$. In (Schwinghammer et al. 2009) using Lemma 5.3 we show that there exists a reduction sequence $T_H(p') \xrightarrow{\text{ev}, \leq n} q''$, where q'' is successful.
- The reduction $T_H(p) \xrightarrow{\text{ev}} q$ is the first reduction of an encoded $\text{HANDLE.NEW}(\text{ev})$ reduction, i.e. it is a $\beta\text{-CBV}(\text{ev})$ -reduction and $p \xrightarrow{\text{HANDLE.NEW}(\text{ev})} p'$. Then again it is possible to construct a reduction sequence $T_H(p') \xrightarrow{\text{ev}, \leq n} q''$ where q'' is successful.

In all cases we can apply the induction hypothesis to $T_H(p')$. $\square$

Proposition 5.6. *The translation T_H is convergence equivalent.*

Proof. Propositions 5.4 and 5.5 show that T_H preserves and reflects may-convergence. Preservation and reflection of must-convergence follows by showing that T_H preserves and reflects may-divergence. The proofs are similar, where the base cases of the inductions are valid since from preservation and reflection of may-convergence it also follows that T_H preserves and reflects must-divergence. $\square$

Compositionality and convergence equivalence of T_H imply:

Theorem 5.7 (Observational correctness and adequacy of T_H). *The translation T_H is observationally correct and adequate.*

Since the identity translation $\iota_H : \lambda^\tau(\text{fcb}) \rightarrow \lambda^\tau(\text{fchb})$ which we introduced above is an embedding and obviously compositional and convergence equivalent, and since the translation T_H is injective on types, Proposition 2.3 is applicable:

Corollary 5.8 (Full abstraction). *The translations T_H and ι_H are fully abstract.*

Using ι_H and ι_B (see Theorem 4.11) we can define direct translations T'_B, T'_H from $\lambda^\tau(\text{fcb})$ into $\lambda^\tau(\text{fch})$ and vice versa, by setting $T'_B = T_B \circ \iota_H$ and $T'_H = T_H \circ \iota_B$. Since full abstraction and adequacy are preserved under composition, we obtain an adequate translation $T'_B : \lambda^\tau(\text{fcb}) \rightarrow \lambda^\tau(\text{fch})$ and a fully abstract translation $T'_H : \lambda^\tau(\text{fch}) \rightarrow \lambda^\tau(\text{fcb})$.

Remark 5.9. *Adequacy of T_H (T'_H , resp.) implies that used handles are equivalent to suspended black holes, i.e. $x \text{ h } \bullet \sim x \xleftarrow{\text{susp}} x$ in $\lambda^\tau(\text{fchb})$ and $\lambda^\tau(\text{fch})$. This follows, since the translations are syntactically equal $\lambda^\tau(\text{fcb})$ processes, i.e. $T_H(x \text{ h } \bullet) \equiv x \xleftarrow{\text{susp}} x \equiv T_H(x \xleftarrow{\text{susp}} x)$.*

$$\begin{aligned}
\text{buffer} &\triangleq \lambda_{-}. \text{let } x_g = \mathbf{ref} \text{ false}, x_p = \mathbf{ref} \text{ true}, x_s = \mathbf{ref} \text{ None} \\
&\quad \text{in thread } \lambda_{-}. \langle x_g, x_p, x_s \rangle \\
\text{get} &\triangleq \mathbf{thread}(\lambda \text{ ig. } \lambda x. \\
&\quad \text{let } \langle x_g, x_p, x_s \rangle = x, \text{getOK} = \mathbf{exch}(x_g, \text{false}) \\
&\quad \text{in if } \text{getOK} \text{ then let } \text{Some } v = \mathbf{exch}(x_s, \text{None}) \\
&\quad \quad \text{dummy} = \mathbf{exch}(x_p, \text{true}) \quad \text{in } v \\
&\quad \text{else ig } x) \\
\text{put} &\triangleq \mathbf{thread}(\lambda \text{ ip. } \lambda \langle x, v \rangle. \\
&\quad \text{let } \langle x_g, x_p, x_s \rangle = x, \text{putOK} = \mathbf{exch}(x_p, \text{false}) \\
&\quad \text{in if } \text{putOK} \text{ then } \mathbf{exch}(x_s, \text{Some } v); \mathbf{exch}(x_g, \text{true}); \mathbf{unit} \\
&\quad \text{else ip } \langle x, v \rangle)
\end{aligned}$$

Figure 11. Busy-wait encoding of buffers

6. Encoding Buffers with Cells and Busy-wait

The results obtained above lead to the question if buffers (and therefore also handled futures) can be encoded in a calculus without either synchronization primitive. In this section we partly answer this question. We encode buffers as cells and buffer operations as operations on cells, by giving a translation $T_{RB} : \lambda^\tau(\text{fcb}) \rightarrow \lambda^\tau(\text{fc})$. Here, the target calculus $\lambda^\tau(\text{fc})$ is like $\lambda^\tau(\text{fcb})$ but without buffers. The translation will be a busy-wait encoding, and our semantics will show that this is an adequate encoding. However, at least from a performance point of view, these kinds of encoding should be avoided, since the knowledge where processes are waiting for a certain event gets lost.

We will use a data type $\text{option}(\tau)$ with the nullary constructor None and unary constructor Some : $\tau \rightarrow \text{option}(\tau)$. This simplifies the encoding since no dummy values are needed. Fig. 11 shows the encoding of the buffer operations that induces the translation $T_{RB} : \lambda^\tau(\text{fcb}) \rightarrow \lambda^\tau(\text{fc})$. On buffers it is defined as follows:

$$\begin{aligned}
T_{RB}(x \text{ b } v) &\triangleq (\nu x_g, x_p, x_s)(x \leftarrow (x_g, x_p, x_s) \mid x_p \text{ c false} \\
&\quad \mid x_g \text{ c true} \mid x_s \text{ c Some } T_{RB}(v)) \\
T_{RB}(x \text{ b } -) &\triangleq (\nu x_g, x_p, x_s)(x \leftarrow (x_g, x_p, x_s) \mid x_p \text{ c true} \\
&\quad \mid x_g \text{ c false} \mid x_s \text{ c None})
\end{aligned}$$

This encoding is compositional and preserves success, i.e., p is successful iff $T_{RB}(p)$ is. Moreover, Theorem 3.4 and adequacy of the identity transformation $\text{id} : \lambda^\tau(\text{fc}) \rightarrow \lambda^\tau(\text{fch})$ show that the reduction rules of $\lambda^\tau(\text{fch})$, except for $\text{CELL.EXCH}(\text{ev})$ and the handle-rules, and the transformations of Fig. 7, are also correct transformations for $\lambda^\tau(\text{fc})$. This allows us to prove convergence equivalence of $T_{RB}(p)$, which implies correctness:

Theorem 6.1 (Observational correctness and adequacy). *The translation $T_{RB} : \lambda^\tau(\text{fcb}) \rightarrow \lambda^\tau(\text{fc})$ is observationally correct and adequate.*

As a corollary, using the translation $T'_H : \lambda^\tau(\text{fch}) \rightarrow \lambda^\tau(\text{fcb})$ described at the end of the previous section, we obtain an observationally correct and adequate translation $T_{RB} \circ T'_H$ from $\lambda^\tau(\text{fch})$ into $\lambda^\tau(\text{fc})$.

Essentially, these translations show that handled futures and buffers are not necessary from a semantic viewpoint. However, the blocking and waiting by queuing is preferable to the busy-wait implementation given in this section, since the machine can keep track of suspended processes and notify them after changes. Accordingly, we obtain the following theoretical challenge, that we currently leave open:

Show that there is an efficiency advantage of the handle-removing and buffer-removing translations T'_H, T'_B over the busy-wait translations.

A proof of this challenge could be possible by extending our current development and arguments of the semantical theory of may- and must-convergence. Specifically, one must find an appropriate notion of corresponding reduction trees, and analyze the lengths of the corresponding reduction sequences more carefully.

Finally, to complete the diagram from the introduction, it remains to relate $\lambda^\tau(\text{fc})$ to the calculus $\lambda(\text{fh})$. We achieve this in two steps: First, the data constructors and case-expressions of $\lambda^\tau(\text{fc})$ are encoded in an (untyped) intermediate calculus that only has concurrent and lazy futures and reference cells. In a second step, this intermediate calculus is embedded into the calculus $\lambda(\text{fh})$ (which additionally features handled futures). The translations used in both steps are observationally correct and adequate, hence this also holds for their composition. The details of this construction can be found in the technical report (Schwinghammer et al. 2009).

7. Discussion and Related Work

In this paper we have proposed a method for specifying and reasoning about implementations, based on semantics-preserving translations. We have proved that concurrent buffers and handled futures are equivalent synchronization primitives in the lambda calculus with futures, in the sense that each can correctly encode the other. This result can be seen as an extended case study of our method, and illustrates how recent proof techniques based on observational semantics permit to prove for a first time the equivalence of various concurrency primitives of realistic concurrent programming languages.

To complete the picture in our particular setting, there are two immediate open questions: can we make precise the intuition that the translations T_H, T_B that encode blocking by waiting and queuing and vice versa, improve upon the busy-wait encoding? ; and is the translation T_B fully abstract?

Questions of expressiveness have been addressed mainly in the pi-calculus and basic process calculi (Parrow 2008; Gorla 2008); we are not aware of previous work on formally relating synchronization primitives in concurrent high-level languages w.r.t. **contextual semantics**. Similar issues, concerning properties of translations, arise in the verification of compilers, an ongoing research topic (e.g. Leroy and Blazy 2008). However, in this context usually only (simpler) simulation properties for closed programs, rather than open program fragments, are established.

One technique for relating different primitives for *sequential* languages are proofs of representation independence, which guarantees that an invariant between two implementations (or an implementation and its specification) is preserved in all programs (Reynolds 1983). While recent work (Ahmed et al. 2009; Bohr and Birkedal 2006; Reddy and Yang 2004) has extended this method from functional to stateful higher-order languages, it is not clear to us whether it can also be adapted to concurrent languages. Our Section 4.4 can be viewed as a result of this kind, which is obtained by analysis of reduction sequences.

Proving similar correctness and expressiveness results for other base languages is possible in principle, but requires to establish a sufficiently rich equational theory first. Here, we derived sufficiently many equivalences for our proofs via adequate translations into “smaller” core calculi. Alternatives to this approach may be bisimulation methods, based on suitable labelled transition systems. For instance, bisimulation methods have been developed for an untyped, stateful higher-order language by Koutavas and Wand (2006), and for fragments of Concurrent ML (e.g., by Jeffrey and Rathke 2004). However, it is open in which way bisimilarity can characterize contextual semantics w.r.t. may- and must-convergence, in particular for languages like the ones considered in this paper.

References

- A. Ahmed, D. Dreyer, and A. Rossberg. State-dependent representation independence. *POPL'09*, pages 340–353. ACM Press, 2009.
- N. Bohr and L. Birkedal. Relational reasoning for recursive types and references. *APLAS'06*, LNCS 4279, pages 79–96. Springer, 2006.
- A. Carayol, D. Hirschkoﬀ, and D. Sangiorgi. On the representation of McCarthy’s amb in the pi-calculus. *TCS*, 330(3):439–473, 2005.
- R. De Nicola and M. Hennessy. Testing equivalences for processes. *TCS*, 34:83–133, 1984.
- C. Fournet, F. Le Fessant, L. Maranget, and A. Schmitt. JoCaml: A language for concurrent distributed and mobile programming. In *Advanced Functional Programming*, LNCS 2638, pages 129–158. Springer, 2002.
- D. Gorla. Towards a unified approach to encodability and separation results for process calculi. *CONCUR '08*, pages 492–507, Springer, 2008.
- A. Jeffrey and J. Rathke. A theory of bisimulation for a fragment of concurrent ML with local names. *TCS*, 323(1-3):1–48, 2004.
- S. Peyton Jones, A. Gordon, and S. Finne. Concurrent Haskell. *POPL'96*, pages 295–308. ACM Press, 1996.
- V. Koutavas and M. Wand. Small bisimulations for reasoning about higher-order imperative programs. *POPL'06*, pages 141–152. ACM, 2006.
- X. Leroy and S. Blazy. Formal verification of a C-like memory model and its uses for verifying program transformations. *J. Autom. Reasoning*, 41(1):1–31, 2008.
- B. Liskov and L. Shriram. Promises: linguistic support for efficient asynchronous procedure calls in distributed systems. *PLDI'88*, pages 260–267. ACM Press, 1988.
- J. Niehren, J. Schwinghammer, and G. Smolka. A concurrent lambda calculus with futures. *TCS* 364(3):338–356, November 2006.
- J. Niehren, D. Sabel, M. Schmidt-Schauß, and J. Schwinghammer. Observational semantics for a concurrent lambda calculus with reference cells and futures. *23rd MFPS*, ENTCS 173, pages 313–337. Elsevier, 2007.
- C.-H. L. Ong. Non-determinism in a functional setting. In *LICS '93*, pages 275–286. IEEE Computer Society, 1993.
- J. Parrow. Expressiveness of process algebras. *ENTCS*, 209:173–186, 2008.
- S. Peyton Jones, A. Gordon, and S. Finne. Concurrent Haskell. *POPL'96*, pages 295–308, ACM Press, 1996.
- U. S. Reddy and H. Yang. Correctness of data representations involving heap data structures. *Sci. Comput. Program.*, 50(1–3):129–160, 2004.
- A. Rensink and W. Vogler. Fair testing. *Inform. and Comput.*, 205(2):125–198, 2007.
- J. H. Reppy. *Concurrent Programming in ML*. Cambridge University Press, 1999.
- J. C. Reynolds. Types, abstraction, and parametric polymorphism. In *Information Processing '83*, pages 513–523. North-Holland, 1983.
- J. G. Riecke. Fully abstract translations between functional languages. *POPL'91*, pages 245–254, ACM-Press, 1991.
- E. Ritter and A. M. Pitts. A fully abstract translation between a lambda-calculus with reference types and standard ml. *TLCA'95*, LNCS 902, Springer, 1995.
- A. Rossberg, D. Le Botlan, G. Tack, T. Brunklaus, and G. Smolka. *Trends in Functional Programming*, volume 5, chapter Alice Through the Looking Glass, pages 79–96. Munich, Germany, February 2006.
- D. Sabel and M. Schmidt-Schauß. A call-by-need lambda calculus with locally bottom-avoiding choice: context lemma and correctness of transformations. *Math. Struct. Comp. Sci.*, 18(3):501–553, 2008.
- M. Schmidt-Schauß, J. Niehren, D. Sabel, and J. Schwinghammer. Adequacy of compositional translations for observational semantics. In *5th IFIP*, IFIP 273, pages 521–535. Springer, 2008.
- M. Schmidt-Schauß and D. Sabel. On generic context lemmas for lambda calculi with sharing. Frank report 27, Inst. f. Informatik, J.W.Goethe-University, Frankfurt, 2007.
- M. Schmidt-Schauß, J. Niehren, J. Schwinghammer, and D. Sabel. Adequacy of compositional translations for observational semantics. Frank report 33, Inst. f. Informatik, J.W.Goethe-University, Frankfurt, 2009.
- J. Schwinghammer, D. Sabel, J. Niehren, and M. Schmidt-Schauß. On correctness of buffer implementations in a concurrent lambda calculus with futures. Frank report 37, Inst. f. Informatik, J.W.Goethe-University, Frankfurt, 2009.