

UNIVERSITY OF OSLO
Department of Informatics

**Behavioral interface
description of an
object-oriented
language with
futures and
promises¹**

Research Report No.
364

Erika Ábrahám

Immo Grabe

Andreas Grüner

Martin Steffen

ISBN 82-7368-322-2

ISSN 0806-3036

October 2007



Behavioral interface description of an object-oriented language with futures and promises[†]

Erika Ábrahám Immo Grabe Andreas Grüner
Martin Steffen

12th October 2007

Abstract

This paper formalizes the observable interface behavior of an concurrent, object-oriented language with futures and promises. The calculus captures the core of *Creol*, a language, featuring in particular asynchronous method calls and, since recently, first-class futures.

The focus of the paper are *open* systems and we formally characterize their behavior in terms of interactions at the interface between the program and its environment. The behavior is given by transitions between typing judgments, where the absent environment is represented abstractly by an assumption context. A particular challenge is the safe treatment of promises: The erroneous situation that a promise is fulfilled twice, i.e., bound to code twice, is prevented by a resource aware type system, enforcing linear use of the write-permission to a promise. We show subject reduction and the soundness of the abstract interface description.

Keywords: concurrent object-oriented languages, Creol, formal semantics, concurrency, futures and promises, open systems, observable behavior

1 Introduction

How to marry concurrency and object-orientation has been a long-standing issue; see e.g., [11] for an early discussion of different design choices. The thread-based model of concurrency, prominently represented by languages like *Java* and *C#*, has been recently criticized, especially in the context of *component-based* software development. As the word indicates, components are (software) artifacts intended for composition, i.e., open systems, interacting with a surrounding environment. To compare different concurrency models for open systems on a solid mathematical basis, a semantical description of the interface behavior is needed, and this is what we provide in this work.

*Part of this work has supported by the NWO/DFG project Mobi-J (RO 1122/9-4) and by the EU-project IST-33826 *Crede: Modeling and analysis of evolutionary structures for distributed services*. For more information, see <http://credo.cwi.nl>

[†]Part of this work has supported by the NWO/DFG project Mobi-J (RO 1122/9-4) and by the EU-project IST-33826 *Crede: Modeling and analysis of evolutionary structures for distributed services*. For more information, see <http://credo.cwi.nl>

We present an *open semantics* for the core of the *Creol* language [25, 43], an object-oriented, concurrent language, featuring in particular asynchronous method calls and, since recently [27], first-class futures.

Futures and promises

A *future*, very generally, represents a result yet to be computed. It acts as a proxy for , or reference to, the delayed result from some piece of code (e.g., a method or a function body in an object-oriented, resp. a functional setting). As the consumer of the result can proceed its own execution until it actually needs it, futures provide a natural, lightweight, and (in a functional setting) transparent mechanism to introduce parallelism into a language. Since their introduction in *Multilisp* [36][13], futures have been used in various languages like Alice ML [45, 9, 58], E [28], the ASP-calculus [18], Creol, and others. A *promise* is a generalization¹ insofar as the reference to the result on the one hand, and the code to calculate the result on the other, are not created at the same time; instead, a promise can be created and only later, after possibly passing it around, be bound to the code (the promise is *fulfilled*).

The notion of futures goes back to functional programming languages. In the functional setting, futures are annotations to side-effect-free expressions², that can be computed in parallel to the rest of the program. If some program code needs the result of a future, its execution blocks until the future’s evaluation is completed and the result value is automatically fetched back (*implicit* futures). An important property of future-based functional programs is, that future annotations do not change the functionality: the observable behavior of an annotated program equals the observable behavior of its non-annotated counterpart.

Interface behavior

An open program interacts with its environment via message exchange. The interface behavior of an open program C can be characterized by the set of all those message sequences (traces) t , for which there *exists* an environment E such that C and E exchange the messages recorded in t . Thereby we abstract away from any concrete environment, but consider only environments that are compliant to the language restrictions (syntax, type system, etc.). Consequently, interactions are not arbitrary traces $C \xrightarrow{t}$; instead we consider behaviors $C \parallel E \xrightarrow[t]{t} \acute{C} \parallel \acute{E}$ where E is an *realizable* environment and \bar{t} is complementary to t . To account for the abstract environment (“there exists an E s.t. ...”), the open semantics is given in an *assumption-commitment* way:

$$\Delta \vdash C : \Theta \xRightarrow{t} \acute{\Delta} \vdash \acute{C} : \acute{\Theta} ,$$

where Δ (as an abstract version of E) contains the *assumptions* about the environment, and dually Θ the *commitments* of the component. Abstracting away also from C gives a language characterization by the set of all possible traces between any component and any environment.

¹The terminology concerning futures, promises, and related constructs is not too consistent in the literature. Sometimes, the two words are used as synonyms. Interested in the observable differences between futures and promises, we distinguish the concepts and thus follow the terminology as used e.g., in λ_{fut} , Alice ML, and the definition given in Wikipedia.

²Though in e.g. *Multilisp* also side-effect expressions can be computed in parallel, but still under the restriction that the observable behavior equals that of the sequential counterpart.

Such a behavioral interface description is relevant and useful for the following reasons. 1) The set of possible traces is more restricted than the one obtained when ignoring the environments. When reasoning about the trace-based behavior of a component, e.g., in compositional verification, with a more precise characterization one can carry out stronger arguments. 2) When using the trace description for *black-box testing*, one can describe test cases in terms of the interface traces and then synthesize appropriate test drivers from it. Clearly, it makes no sense to specify impossible interface behavior, as in this case one cannot generate a corresponding tester. 3) A representation-independent behavior of open programs paves the way for a compositional semantics and allows furthermore optimization of components: only if two components show the same external behavior, one can replace one for the other without changing the interaction with any environment. 4) The formulation gives insight into the semantical nature of the language, here, the observable consequences of futures and promises. This helps to compare alternatives, e.g., the Creol concurrency model with *Java*-like threading.

Results

The paper formalizes the abstract interface behavior for concurrent object-oriented languages with futures and promises. The contributions are the following.

Concurrent object calculus with futures and promises We formalize a class-based concurrent language featuring futures and promises. The formalization is given as a typed, imperative object calculus in the style of [1] resp. one of its concurrent extensions. The operational semantics for components distinguishes unobservable component-internal steps from external steps which represent observable component-environment interactions. We present the semantics in a way that facilitates comparison with *Java*'s multi-threading concurrency model, i.e., the operational semantics is formulated so that the multi-threaded concurrency as (for instance) in *Java* and the one here based on futures are represented similarly.

Linear type system for promises The calculus extends the semantic basis of *Creol* as given for example in [27] with promises. Promises can refer to a computation with code bound to it later, where the binding is done at most once. To guarantee such a *write-once* policy when passing around promises, we refine the type system introducing two type constructors

$$[T]^{+-} \quad \text{and} \quad [T]^+$$

representing a reference to a promise that can still be written (and read), with result type T , resp. that has a *read*-permission. The write permission constitutes a resource which is consumed when the promise is fulfilled. The resource-aware type system is therefore formulated in a *linear* manner wrt. the write permissions and resembles in intention the one in [53] for a functional calculus with references. Our work is more general, in that it tackles the problem in an object-oriented setting (which, however, conceptually does not pose much complications), and in that we do not consider closed systems, but open components. Also this aspect of openness is not dealt with in [27]. Additionally, the type system presented here is simpler as in [53], as it avoids the representation of the promise-concept by so-called *handled futures*.

Soundness of the abstractions We show soundness of the abstractions, which includes

C	$::= \mathbf{0} \mid C \parallel C \mid \underline{v(n:T).C} \mid n[(O)] \mid \underline{n[n, F, L]} \mid \underline{n\langle t \rangle}$	program
O	$::= F, M$	object
M	$::= l = m, \dots, l = m$	method suite
F	$::= l = f, \dots, l = f$	fields
m	$::= \zeta(n:T).\lambda(x:T, \dots, x:T).t$	method
f	$::= \zeta(n:T).\lambda().v \mid \zeta(n:T).\lambda().\perp_{n'}$	field
t	$::= v \mid \text{stop} \mid \text{let } x:T = e \text{ in } t$	thread
e	$::= t \mid \text{if } v = v \text{ then } e \text{ else } e \mid \text{if } \text{undef}(v.l()) \text{ then } e \text{ else } e$ $\mid \text{promise } T \mid \text{bind } n.l(\vec{v}) : T \leftrightarrow n \mid \underline{\text{set } v \mapsto n} \mid v.l() \mid v.l := \zeta(s:n).\lambda().v$ $\mid \text{new } n \mid \underline{\text{claim}@}(n, n) \mid \underline{\text{get}@}n \mid \underline{\text{suspend}(n)} \mid \underline{\text{grab}(n)} \mid \underline{\text{release}(n)}$	expr.
v	$::= x \mid n \mid ()$	values
L	$::= \perp \mid \top$	lock status

Table 1: Abstract syntax

- *subject reduction*, i.e., preservation of well-typedness under reduction. Subject reduction is not just proven for a closed system (as usual), but for an open system interacting with its environment. Subject reduction implies
- *absence of run-time errors* like “message-not-understood”, also for open systems.
- *soundness* of the interface behavior characterization, i.e., all possible interaction behavior is included in the abstract interface behavior description.
- for promises: absence of *write-errors*, i.e. the attempt to fulfill a promise twice.

The paper is organized as follows. Section 2 defines the syntax, the type system, and the operational semantics, split into an internal one and one for open systems. Section 3 describes the interface behavior. Section 4 concludes with related and future work. For more details and for the proofs see [2].

2 Calculus

This section presents the calculus, based on a version of the *Creol*-language with first-class futures [27] and extended with promises. It is a concurrent variant of an imperative, object-calculus in the style of the ones from [1].

2.1 Syntax

The abstract syntax in Table 1 distinguishes between *user* syntax and *run-time* syntax (the latter is underlined). The user syntax contains the phrases in which programs are written; the run-time syntax contains syntactic constituents additionally needed to express the executing program in the operational semantics.

Names n refer to classes, objects, threads, and to references to futures and promises. We use o and its syntactic variants for objects and c for classes, and n when being un-specific. The unit value is represented by $()$. A component C is a collection of classes, objects, and threads, with $\mathbf{0}$ being the empty component. A class $c[(O)]$ carries a name c

and defines its methods and fields in O . A method $\zeta(s:c).\lambda(\vec{x}:\vec{T}).t$ provides the method body t abstracted over the ζ -bound “self” parameter s and the formal parameters \vec{x} . For uniformity, fields are represented as methods without parameters (except self), with a body being either a value or yet undefined. An object $o[c, F, L]$ with identity o keeps a reference to the class c it instantiates, stores the current value F of its fields, and maintains a *binary lock* L indicating whether any code is currently active inside the object (in which case the lock is taken) or not (in which case the lock is free). The symbols \top , resp., \perp , indicate that the lock is taken, resp., free. Note that the methods are stored in the classes but the fields are kept in the objects, of course. In freshly created objects, the lock is free, and all fields carry the undefined reference \perp_c , where class name c is the (return) type of the field.

Besides objects and classes, the dynamic configuration of a program contains incarnations of method bodies, written $n\langle t \rangle$, as active entities. The term t is basically a sequence of expressions, where the let-construct is used for sequencing and for local declarations.³ During execution, $n\langle \text{let } x:T = t \text{ in } x \rangle$ contains in t the currently running code of a method body and the result will be stored in the local variable x . When evaluated, the thread is of the form $n'\langle \text{set } v \mapsto n \rangle$ and the value can be accessed via n , the future reference, or future for short, where $\text{set } v \mapsto n$ is an auxiliary expression.⁴

We use f for instance variables or fields and $l = \zeta(s:T).\lambda().v$, resp. $l = \zeta(s:T).\lambda().\perp_c$ for field variable definition. Field access is written as $v.l()$ and field update as $v'.l := \zeta(s:T).\lambda().v$. By convention, we abbreviate the latter constructs by $l = v$, $l = \perp_c$, $v.l$, and $v'.l := v$. We will also use v_\perp to denote either a value v or a symbol \perp_c for being undefined. Note that the syntax does not allow to set a field back to undefined. Direct access (read or write) to fields across object boundaries is forbidden, and we do not allow method update. Instantiation of a new object from class c is denoted by $\text{new } c$.

Expressions include especially *promise* T for creating a new promise, and $\text{bind } o.l(\vec{v}) : T \hookrightarrow n$ for binding the method call $o.l(\vec{v})$ with return type T to promise n . Asynchronous method calls, central to *Creol*'s concurrency model, are a derived concept. An asynchronous call, written $o@l(\vec{v})$ is syntactic sugar for creating a new promise and immediately binding $o.l(\vec{v})$ to it. Further, the expressions *claim*, *get*, *suspend*, *grab*, and *release* deal with communication and synchronization. The expression $\text{claim}@ (n, o)$ is the attempt to obtain the result of a method call from the future named n while in possession of the lock of object o . Executing $\text{release}(o)$ relinquishes the lock of the object o , giving other threads the chance to be executed in its stead, when succeeding to grab the lock via $\text{grab}(o)$. Executing $\text{suspend}(o)$ causes the activity to relinquish and re-grab the lock of object o (see the operational rules in Section 2.3.1 below). We assume by convention, that when appearing in methods of classes, the claim- and the suspend-command only refer to the self-parameter *self*, i.e., they are written $\text{claim}@ (n, \text{self})$ and $\text{suspend}(\text{self})$.⁵

2.2 Type system

The calculus is typed and the available types are given in the following grammar:

³ $t_1; t_2$ (sequential composition) abbreviates $\text{let } x:T = t_1 \text{ in } t_2$, where x does not occur free in t_2 .

⁴The reason why an evaluated future n is represented by $n'\langle \text{set } v \mapsto n \rangle$ and not by $n\langle v \rangle$, which might look more natural, is technical. In the operational semantics, the reference n' is hidden. Technically, the representation allows to achieve subject reduction for the open semantics, without exposing the status of the future n .

⁵For the run-time constructs *grab* and *release*, we need not impose the analogous restriction, as it is guaranteed by the operational semantics.

$$\begin{aligned}
T & ::= B \mid \text{Unit} \mid [T]^{+-} \mid [T]^+ \mid [!U, \dots, !U] \mid [!(U, \dots, !U)] \mid n \\
U & ::= T \times \dots \times T \rightarrow T
\end{aligned}$$

Besides base types B (left unspecified), Unit is the type of the unit value $()$. Types $[T]^{+-}$ and $[T]^+$ represent the reference to a future which will return a value of type T , in case it eventually terminates. $[T]^{+-}$ indicates that the promise has not yet been fulfilled, i.e., it represents the write-permission to a promise (which implies read-permission at the same time). $[T]^+$ represents read-only permission to a future. The read/write capability is more specific than read-only, which is expressed by the (rather trivial) subtyping relation generated by $[T]^{+-} \leq [T]^+$, accompanied by the usual subsumption rule. Furthermore, $[-]^+$ acts monotonely, and $[-]^{+-}$ invariantly wrt. subtyping. When not interested in the access permission, we just write $[T]$.

The name of a class serves as the type for the named instances of the class. We need as auxiliary type construction the type or interface of unnamed objects, written $[l_1:U_1, \dots, l_k:U_k]$ and the interface type for classes, written $[(l_1:U_1, \dots, l_k:U_k)]$. We allow ourselves to write \vec{T} for $T_1 \times \dots \times T_k$ etc. where we assume that the number of arguments match in the rules, and write $\text{Unit} \rightarrow T$ for $T_1 \times \dots \times T_k \rightarrow T$ when $k = 0$.

We are interested in the behavior of well-typed programs, only, and the section presents the type system to characterize those. As the operational rules later, the derivation rules for typing are grouped into two sets: one for typing on the level of components, i.e., global configurations, and secondly one for their syntactic sub-constituents.

Table 2 defines the typing on the level of configurations, i.e., for “sets” of objects, classes, and threads. On this level, the typing judgments are of the form

$$\Delta \vdash C : \Theta, \tag{1}$$

where Δ and Θ are *name contexts*, i.e., finite mappings from names to types. In the judgment, Δ plays the role of the typing assumptions about the environment, and Θ of the commitments of the configuration, i.e., the names offered to the environment. Sometimes, the words required and provided interface are used to describe their dual roles. Δ must contain at least all external names referenced by C and dually Θ mentions the names offered by C , which constitute the static interface information. A pair Δ and Θ of assumption and commitment context with disjoint domains are called *well-formed*.

$\frac{}{\Delta \vdash \mathbf{0} : ()}$ T-EMPTY	$\frac{\Delta, \Theta_2 \vdash C_1 : \Theta_1 \quad \Delta, \Theta_1 \vdash C_2 : \Theta_2}{\Delta \vdash C_1 \parallel C_2 : \Theta_1, \Theta_2}$ T-PAR	$\frac{\Delta \vdash C : \Theta, n:T}{\Delta \vdash \nu(n:T).C : \Theta}$ T-NU
$\frac{; \Delta, c:T \vdash [(O)] : T}{\Delta \vdash c[(O)] : (c:T)}$ T-NCLASS	$\frac{; \Delta \vdash c : [(T_F, T_M)] \quad ; \Delta, o:c \vdash [F] : [T_F]}{\Delta \vdash o[c, F, l] : (o:c)}$ T-NOBJ	
$\frac{; [\Delta], n:[T]^+ \vdash t : T}{\Delta \vdash n(t) : (n:[T]^+)}$ T-NTHREAD	$\frac{\Delta' \leq \Delta \quad \Theta \leq \Theta' \quad \Delta \vdash C : \Theta}{\Delta' \vdash C : \Theta'}$ T-SUB	

Table 2: Typing (components)

The empty configuration $\mathbf{0}$ is well-typed in any context and exports no names (cf. rule T-EMPTY). Two configurations in parallel can refer mutually to each other's commitments, and together offer the (disjoint) union of their names (cf. rule T-PAR). It will be an invariant of the operational semantics that the identities of parallel entities are disjoint wrt. the mentioned names. Therefore, Θ_1 and Θ_2 in the rule for parallel composition are merged disjointly, indicated by writing Θ_1, Θ_2 (analogously for the assumption contexts). In combination with the rest of the rules (in particular T-BIND below), this assures that a promise cannot be fulfilled by the component and the environment at the same time. The ν -binder hides the bound object or the name of the future inside the component (cf. rule T-NU). In the T-NU-rule, we assume that the bound name n is new to Δ and Θ . Let-bound variables are *stack* allocated and checked in a stack-organized variable context Γ . Names created by `new` are *heap* allocated and thus checked in a "parallel" context (cf. again the assumption-commitment rule T-PAR). The rules for named classes introduce the name of the class and its type into the commitment (cf. T-NCLASS). The code of the class $\llbracket O \rrbracket$ is checked in an assumption context where the name of the class is available. An instantiated object will be available in the exported context Θ by rule T-NOBJ. Named threads $n\langle t \rangle$ are treated by rule T-NTHREAD, where the type $[T]^+$ of the future reference n is matched against the result type T of thread t . As obviously future n is already fulfilled in $n\langle t \rangle$, its type exports read-permission, only. For a named thread $n\langle t \rangle$ in rule T-NTHREAD to be well-typed, the code t is checked using the assumptions Δ of the conclusion but *without* using write-permissions mentioned in Δ , expressed by $[\Delta]$. On types, the $[_]$ operation is defined as $[[T]^{+-}] = [T]^+$ and as identity on all other types. The definition is lifted pointwise to binding contexts. The last rule is a rule of subsumption, expressing a simple form of subtyping: we allow that an object respectively a class contains *at least* the members which are required by the interface. This corresponds to width subtyping. Note, however, that each named object has exactly one type, namely its class.

Definition 2.1 (Subtyping). *The relation \leq on types is defined as identity for all types except for $[T]^{+-} \leq [T]^+$ (mentioned above) and object interfaces, where we have:*

$$\llbracket l_1:U_1, \dots, l_k:U_k, l_{k+1}:U_{k+1}, \dots \rrbracket \leq \llbracket l_1:U_1, \dots, l_k:U_k \rrbracket.$$

For well-formed name contexts Δ_1 and Δ_2 , we define in abuse of notation $\Delta_1 \leq \Delta_2$, if Δ_1 and Δ_2 have the same domain and additionally $\Delta_1(n) \leq \Delta_2(n)$ for all names n .

The same definition is applied, of course, also for name contexts Θ , used for the commitments. The relations \leq are obviously reflexive, transitive, and antisymmetric.

Next we formalize the typing for objects and threads and their syntactic sub-constituents. Especially the treatment of the write-permission requires care: The capability to write to a promise is consumed by the bind-operation as it should be done only once. This is captured by a *linear* type system where the execution of a thread or an expression may change the involved types. The judgments are of the form

$$\Gamma; \Delta \vdash e : T :: \acute{\Gamma}, \acute{\Delta}, \quad (2)$$

where the change from Γ and Δ to $\acute{\Gamma}$ and $\acute{\Delta}$ reflects the potential consumption of write permissions when executing e . The consumption is only potential, as the type system statically overapproximates the run-time behavior, of course. The typing is given in Tables 3 and 4. For brevity, we write $\Delta; \Gamma \vdash e : T$ for $\Delta; \Gamma \vdash e : T :: \acute{\Gamma}, \acute{\Delta}$, when $\acute{\Gamma} = \Gamma$ and $\acute{\Delta} = \Delta$. Besides assumptions about the provided names of the environment kept in Δ , the typing is done relative to assumptions about occurring free variables.

They are kept separately in a variable context Γ , a finite mapping from variables to types. Apart from the technicalities, treating the write capabilities in a linear fashion is straightforward: one must assure that the corresponding capability is available at most once in the program and is not duplicated when passed around. A promise is no longer available for writing when bound to a variable using the let-construct, or when handed over as argument to a method call or a return.

$\frac{\Gamma; \Delta \vdash c : \llbracket l_1 : U_1, \dots, l_k : U_k \rrbracket \quad \Gamma; \Delta \vdash m_i : U_i :: \Delta \quad m_i = \zeta(s_i; c). \lambda(x_i; T_i). t_i}{\Gamma; \Delta \vdash \llbracket l_1 = m_1, \dots, l_k = m_k \rrbracket : c} \text{T-CLASS}$
$\frac{\Gamma; \Delta \vdash c : \llbracket l_1 : U_1, \dots, l_k : U_k \rrbracket \quad \Gamma; \Delta \vdash f_i : U_i \quad f_i = \zeta(s_i; c). \lambda(). v_{\perp}}{\Gamma; \Delta \vdash \llbracket l_1 = f_1, \dots, l_k = f_k \rrbracket : c} \text{T-OBJ}$
$\frac{\Gamma, \vec{x}; \vec{T}; \Delta, s; c \vdash t : T' :: \vec{\Delta} \quad \Gamma; \Delta \vdash c : T \quad T = \llbracket \dots, l; \vec{T} \rightarrow T', \dots \rrbracket}{\Gamma; \Delta \vdash \zeta(s; c). \lambda(\vec{x}; \vec{T}). t : T.l} \text{T-MEMB}$
$\frac{\Gamma; \Delta, s; c \vdash c : \llbracket \dots, l : \text{Unit} \rightarrow c', \dots \rrbracket}{\Gamma; \Delta \vdash \zeta(s; c). \lambda(). \perp_{c'} : c'} \text{T-UNDEF}$
$\frac{\Gamma; \Delta \vdash v : c \quad \Gamma; \Delta \vdash c : T \quad \Gamma; \Delta \vdash v' : T.l}{\Gamma; \Delta \vdash v.l := v' : c} \text{T-FUPDATE} \quad \frac{\Gamma; \Delta \vdash c : \llbracket T \rrbracket}{\Gamma; \Delta \vdash \text{new } c : c} \text{T-NEWC}$
$\frac{\Gamma_1; \Delta_1 \vdash e : T_1 :: \Gamma_2; \Delta_2 \quad \Gamma_2, x; T_1; \Delta_2 \vdash t : T_2 :: \Gamma_3; \Delta_3}{\Gamma_1; \Delta_1 \vdash \text{let } x : T_1 = e \text{ in } t : T_2 :: \Gamma_3; \Delta_3} \text{T-LET}$
$\frac{\Gamma_1; \Delta_1 \vdash v_1 : T_1 \quad \Gamma_1; \Delta_1 \vdash v_2 : T_1 \quad \Gamma_1; \Delta_1 \vdash e_1 : T_2 :: \Gamma_2; \Delta_2 \quad \Gamma_2; \Delta_1 \vdash e_2 : T_2 :: \Gamma_2; \Delta_2}{\Gamma_1; \Delta_1 \vdash \text{if } v_1 = v_2 \text{ then } e_1 \text{ else } e_2 : T_2 :: \Gamma_2; \Delta_2} \text{T-COND}$
$\frac{\Gamma_1; \Delta_1 \vdash v : c \quad \Gamma_1; \Delta_1 \vdash c : \llbracket \dots, l; \text{Unit} \rightarrow T, \dots \rrbracket \quad \Gamma_1; \Delta_1 \vdash e_1 : T_2 :: \Gamma_2; \Delta_2 \quad \Gamma_1; \Delta_1 \vdash e_2 : T_2 :: \Gamma_2; \Delta_2}{\Gamma; \Delta_1 \vdash \text{if } \text{undef}(v.l()) \text{ then } e_1 \text{ else } e_2 : T_2 :: \Gamma_2; \Delta_2} \text{T-COND}_{\perp}$
$\frac{}{\Gamma; \Delta \vdash \text{stop} : T} \text{T-STOP} \quad \frac{}{\Gamma; \Delta \vdash () : \text{Unit}} \text{T-UNIT} \quad \frac{}{\Gamma; \Delta \vdash \text{set } v \mapsto n : \text{Unit}} \text{T-COMPL}$

Table 3: Typing

Classes, objects, and methods resp. fields have no effect on Δ (see rules T-CLASS, T-OBJ, T-MEMB, and T-UNDEF). Note that especially in T-MEMB, the name context Δ does not change. This does *not* mean, that a method cannot have a side-effect by fulfilling promises, but they are not part of the check of the method *declaration* here. Rule T-CLASS is the introduction rule for class types, the rule of instantiation of a class T-NEWC requires reference to a class-typed name. In the rules T-MEMB and T-FUPDATE we use the meta-mathematical notation $T.l$ to pick the type in T associated with label l , i.e., $T.l$ denotes U , when $T = [\dots, l; U, \dots]$ and analogously for $T = \llbracket \dots, l; U, \dots \rrbracket$. Rules T-CLASS and T-OBJ check the definition of classes resp., of objects against the respective interface type. Note that the type of the self-parameter must be identical to the name of the class, the method resides in. The premises of rule T-MEMB checks the method body in the context Γ appropriately extended with the formal parameters x_i ,

$\frac{}{\Gamma; \Delta \vdash \text{promise } T : [T]^{+-}} \text{T-PROM}$		
$\frac{\Gamma; \Delta \vdash n : [T]^+}{\Gamma; \Delta \vdash \text{claim}@n : T} \text{T-CLAIM}$	$\frac{}{\Gamma; \Delta \vdash \text{get}@n : T} \text{T-GET}$	$\frac{}{\Gamma; \Delta \vdash \text{set } v \mapsto n_1 : \text{Unit}} \text{T-COMPL}$
$\frac{\Gamma(x) = T \quad \hat{\Gamma} = \Gamma \setminus x : T}{\Gamma; \Delta \vdash x : T :: \hat{\Gamma}; \Delta} \text{T-VAR}$	$\frac{\Delta(x) = T \quad \hat{\Delta} = \Delta \setminus n : T}{\Gamma; \Delta \vdash n : T :: \Gamma; \hat{\Delta}} \text{T-NAME}$	
$\frac{\Gamma; \Delta, n; [T]^+ \vdash o : c \quad \Gamma; \Delta, n; [T]^+ \vdash c : [\dots, l: \vec{T} \rightarrow T, \dots]}{\Gamma; \Delta, n : [T]^{+-} \vdash \text{bind } o.l(\vec{v}) : T \hookrightarrow n : [T]^+ :: \hat{\Gamma}; \hat{\Delta}, n; [T]^+} \text{T-BIND}$		
$\frac{\Delta \vdash o : c}{\Gamma; \Delta \vdash \text{suspend}(o) : \text{Unit}} \text{T-SUSPEND}$	$\frac{}{\Gamma; \Delta \vdash \text{grab}(o) : \text{Unit}} \text{T-GRAB}$	$\frac{}{\Gamma; \Delta \vdash \text{release}(o) : \text{Unit}} \text{T-RELEASE}$
$\frac{T \leq T'}{\Gamma_1; \Delta_1 \vdash t : T' :: \Gamma_2; \Delta_2} \text{T-SUB}$		

Table 4: Typing

resp. the context Δ extended by the ζ -bound self-parameter. T-UNDEF works similarly treating the case of an uninitialized field. The terminated expression `stop` and the unit value do not change the capabilities (cf. rules T-STOP and T-UNIT). Note that `stop` has any type (cf. rule T-STOP) reflecting the fact that control never reaches the point *after* `stop`. Further constructs without side effects are the three expressions to manipulate the monitor locks (suspension, lock grabbing, and lock release), object instantiation (T-NEWC), and field update. Wrt. field update in rule T-FUPDATE, the reason why the update has no effect on the contexts is that we do not allow fields to carry a type of the form $[T]^{+-}$. This effectively prevents the passing around of write-permissions via fields. The rule T-LET for let-bindings introduces a local scope. The change from Δ_1 to Δ_2 and further from Δ_2 to Δ_3 (and analogously for the Γ 's) reflects the sequential evaluation strategy: first e is evaluated and afterwards t . For conditionals, both branches must agree on their pre- and post Δ -contexts, which typically means, over-approximating the effect by taking the upper bound on both as combined effect. Note that the comparison of the values in T-COND resp. the check for definedness in T-COND $_{\perp}$ has no side-effect on the contexts. The rule for testing for definedness using `undef` (not shown) works analogously.

Table 4 deals with futures, promises, and especially the linear aspect of consuming and transmitting the write-permissions. The `claim`-command fetches the result value from a future; hence, if the reference n is of type $[T]^+$, the value itself carries type T (cf. rule T-CLAIM). The rule T-GET for `get` works analogously.

The expression `promise` T creates a new promise, which can be read or written and is therefore of type $[T]^{+-}$. Note, however, that the context Δ does *not* change. The reason is that the new name created by `promise` is hidden by a ν -binder immediately after creation and thus does not immediately extend the Δ -context (see the reduction rule PROM below). The binding of a thread t to a promise n is well-typed if the type of n still allows the promise to be fulfilled, i.e., n is typed by $[T]^{+-}$ and not just $[T]^+$. The auxiliary expression `set` $\nu \mapsto n$ is evaluated for its side-effect, only, and is of type

Unit (cf. rule T-COMPL). *claim* dereferences a future, i.e., it fetches a value of type T from the reference of type $[T]^+$. Otherwise, the expression has no effect on Δ , as reading can be done arbitrarily many times. As an aside: in rule T-CLAIM, the type of o is not checked, as by convention, the claim-statement must be used in the form *claim@self* in the user syntax, where *self* is the self-parameter of the surrounding methods. Reduction then preserves well-typedness so a re-check here is not needed. Similar remarks apply to the remaining. The treatment of *get* is analogous (cf. rules T-CLAIM and T-GET). For T-BIND, handing over a promise with read/write permissions as an actual parameter of a method call, the caller loses the right to fulfill the promise. Of course, the caller can only pass the promise to a method which assumes read/write permissions, if itself has the write permission. The loss of the write-permission is specified by setting $\hat{\Delta}$ and $\hat{\Gamma}$ to $\Delta \setminus \vec{v} : \vec{T}$ resp. to $\Gamma \setminus \vec{v} : \vec{T}$. The *difference*-operator $\Delta \setminus n : [T]^{+-}$ removes the *write*-permission for n from the context Δ . In T-BIND, the premise $\Delta; \Gamma, n:[T]^+ \vdash \vec{v} : \vec{T}$ abbreviates the following: assume $\vec{v} = v_1, \dots, v_n$ and $\vec{T} = T_1 \dots T_n$ and let Ξ_1 abbreviate $\Gamma; \Delta, n:[T]^+$. Then $\Xi \vdash \vec{v} : \vec{T}$ means: $\Xi_i \vdash v_i : T_i$ and $\Xi_{i+1} = \Xi_i \setminus T_i$, for all $1 \leq i \leq n$. Note that checking the type of the callee has not side-effect on the bindings. Mentioning a variable or a name removes the write permission (if present) from the respective binding context (cf. T-VAR and T-NAME). The next three rules T-SUSPEND, T-GRAB, and T-RELEASE deal with the expressions for coordination and lock handling; they are typed by Unit. The last rule T-SUB is the standard rule of subsumption.

2.3 Operational semantics

The operational semantics is given in two stages, component internal steps and external ones, where the latter describe the interaction at the interface. Section 2.3.1 starts with component-internal steps, i.e., those definable without reference to the environment. In particular, the steps have no observable external effect. The external steps, presented afterwards in Section 2.3.2, define the interaction between component and environment. They are defined in reference to assumption and commitment contexts. The static part of the contexts corresponds to the static type system from Section 2.2 on component level and takes care that, e.g., only well-typed values are received from the environment.

2.3.1 Internal steps

The internal semantics describes the operational behavior of a *closed* system, not interacting with its environment. The corresponding reduction steps are shown in Table 5, distinguishing between confluent steps \rightsquigarrow and other internal transitions $\xrightarrow{\tau}$, both invisible at the interface. The \rightsquigarrow -steps, on the one hand, do not access the instance state of the objects. They are free of imperative side effects and thus confluent. The $\xrightarrow{\tau}$ -steps, in contrast, access the instance state, either by reading or by writing it, and thus may lead to race conditions. In other words, this part of the reduction relation is in general not confluent.

The first seven rules deal with the basic sequential constructs, all as \rightsquigarrow -steps. The basic evaluation mechanism is substitution (cf. rule RED). Note that the rule requires that the leading let-bound variable is replaced only by *values* v . The operational behavior of the two forms of conditionals are axiomatized by the four COND-rules. Depending on the result of the comparison in the first pair of rules, resp., the result of checking for definedness in the second pair, either the then- or the else-branch is taken.




Figure 1: Claiming a future

In COND_2 , we assume that v_1 does not equal v_2 , as side condition. Evaluating stop terminates the future for good, i.e., the rest of the thread will never be executed as there is no reduction rule for the future $n(\text{stop})$ (cf. rule STOP). The rule FLOOKUP deals with field look-up, where $F'.l(o)()$ stands for $\perp_c[o/s] = \perp_c$, resp., for $v[o/s]$, where $[c, F'] = [c, \dots, l = \zeta(s:c).\lambda().\perp_c, \dots, L]$, if the field is yet undefined, resp., $[c, F'] = [c, \dots, l = \zeta(s:c).\lambda().v, \dots, L]$. In FUPDATE , the meta-mathematical notation $F.l := v$ stands for $(\dots, l = v, \dots)$, when $F = (\dots, l = v', \dots)$. There will be no external variant of the rule for field look-up later in the semantics of open systems, as we do not allow field access across component boundaries. The same restriction holds for field update in rule FUPDATE . A new object as instance of a given class is created by rule $\text{NEW}O_i$. Note that initially, the lock is free and there is not activity associated with the object, i.e., the object is initially passive.

The expression *promise* T creates a fresh promise n . Note that no new thread is yet allocated, as so far nothing more than the name is known. The rule PROM mentions the types T and T' . The typing system assures that the type T is of the form $[S]^{+-}$ for some type S . A promise is fulfilled by the *bind*-command (cf. rule BIND_i), in that the new thread n is put together with the code t_1 to be executed and run in parallel with the rest as $n'(\text{let } x : T = t_1 \text{ in set } x \mapsto n)$ (where n' is hidden). Upon termination, the result is available via the *claim*- and the *get*-syntax (cf. the CLAIM -rules and rule GET_i), but not before the lock of the object is given back again using *release*(o) (cf. rule RELEASE). If the thread is not yet terminated, the requesting thread suspends itself, thereby giving up the lock. The behavior of *claim* is sketched in Figure 1. Note the types of the involved let-bound variables: the future reference is typed by $[T]$, indicating that the value for x will not directly be available, but must be dereferenced first via *claim*. When it comes to claim a future, we added as auxiliary syntax $\text{set } v \mapsto n$. The expression presents an evaluated thread, just in front of the step where the value v is about to be put into the thread named n . The reasons for that additional syntax are largely technical, namely to achieve a clean separation of internal and externally visible behavior, in particular, to get a proper formulation of the subject reduction results. This additional expression requires, that the rules CLAIM_i^1 , CLAIM_i^2 , and GET_1 , dealing with evaluated threads of the form $n\langle v \rangle$, are complemented by the rules CLAIM_i^3 and CLAIM_i^4 , resp. GET_i^2 .

The two operations *grab* and *release* take, resp., give back an object's lock. They are not part of the user syntax, i.e., the programmer cannot directly manipulate the monitor lock. The user can release the lock using the *suspend*-command or by trying to get back the result from a call using *claim*.

The above reduction relations are used modulo *structural congruence*, which captures the algebraic properties of parallel composition and the hiding operator. The basic

axioms for \equiv are shown in Table 6 where in the fourth axiom, n does not occur free in C_1 . The congruence relation is imported into the reduction relations in Table 7. Note that all syntactic entities are always tacitly understood modulo α -conversion.

$n\langle \text{let } x:T = v \text{ in } t \rangle \rightsquigarrow n\langle t[v/x] \rangle$	RED	
$n\langle \text{let } x_2:T_2 = (\text{let } x_1:T_1 = e_1 \text{ in } e) \text{ in } t \rangle \rightsquigarrow n\langle \text{let } x_1:T_1 = e_1 \text{ in } (\text{let } x_2:T_2 = e \text{ in } t) \rangle$	LET	
$n\langle \text{let } x:T = (\text{if } v = v \text{ then } e_1 \text{ else } e_2) \text{ in } t \rangle \rightsquigarrow n\langle \text{let } x:T = e_1 \text{ in } t \rangle$	COND ₁	
$n\langle \text{let } x:T = (\text{if } v_1 = v_2 \text{ then } e_1 \text{ else } e_2) \text{ in } t \rangle \rightsquigarrow n\langle \text{let } x:T = e_2 \text{ in } t \rangle$	COND ₂	
$n\langle \text{let } x:T = (\text{if } \text{undef}(\perp_{\mathcal{C}}) \text{ then } e_1 \text{ else } e_2) \text{ in } t \rangle \rightsquigarrow n\langle \text{let } x:T = e_1 \text{ in } t \rangle$	COND ₁ [⊥]	
$n\langle \text{let } x:T = (\text{if } \text{undef}(v) \text{ then } e_1 \text{ else } e_2) \text{ in } t \rangle \rightsquigarrow n\langle \text{let } x:T = e_2 \text{ in } t \rangle$	COND ₂ [⊥]	
$n\langle \text{let } x:T = \text{stop in } t \rangle \rightsquigarrow n\langle \text{stop} \rangle$	STOP	
$o[c, F', L] \parallel n\langle \text{let } x:T = o.l() \text{ in } t \rangle \xrightarrow{\tau} o[c, F', L] \parallel n\langle \text{let } x:T = F'.l(o)() \text{ in } t \rangle$	FLOOKUP	
$o[c, F, L] \parallel n\langle \text{let } x:T = o.l := v \text{ in } t \rangle \xrightarrow{\tau} o[c, F.l := v, n'] \parallel n\langle \text{let } x:T = o \text{ in } t \rangle$	FUPDATE	
$c[[F, M]] \parallel n\langle \text{let } x:c = \text{new } c \text{ in } t \rangle \rightsquigarrow$		
$c[[F, M]] \parallel v(o:c).(o[c, F, \perp] \parallel n\langle \text{let } x:c = o \text{ in } t \rangle)$	NEW _O _i	
$n'\langle \text{let } x:T' = \text{promise } T \text{ in } t \rangle \rightsquigarrow v(n:T').(n'\langle \text{let } x:T' = n \text{ in } t \rangle)$	PROM	
$c[[F', M]] \parallel o[c, F, l] \parallel n_1\langle \text{let } x:T = \text{bind } o.l(\vec{v}) : T_2 \hookrightarrow n_2 \text{ in } t_1 \rangle \xrightarrow{\tau}$		
$c[[F', M]] \parallel o[c, F, l] \parallel n_1\langle \text{let } x:T = n_2 \text{ in } t_1 \rangle$	BIND _i	
$\parallel v(n':\text{Unit}).(n'\langle \text{let } x:T_2 = \text{grab}(o); M.l(o)(\vec{v}) \text{ in } \text{release}(o); \text{set } x \mapsto n_2 \rangle)$		
$n'\langle \text{set } v \mapsto n_1 \rangle \parallel n_2\langle \text{let } x : T = \text{claim}@ (n_1, o) \text{ in } t \rangle \rightsquigarrow$		
$n'\langle \text{set } v \mapsto n_1 \rangle \parallel n_2\langle \text{let } x : T = v \text{ in } t \rangle$	CLAIM _i ¹	
$\frac{t_2 \neq v}{n'\langle \text{set } t_2 \mapsto n_2 \rangle \parallel n_1\langle \text{let } x : T = \text{claim}@ (n_2, o) \text{ in } t'_1 \rangle \rightsquigarrow}$	CLAIM _i ²	
$n'\langle \text{set } t_2 \mapsto n_2 \rangle \parallel n_1\langle \text{let } x : T = \text{release}(o); \text{get}@n_2 \text{ in } \text{grab}(o); t'_1 \rangle$		
$n_1\langle v \rangle \parallel n_2\langle \text{let } x : T = \text{claim}@ (n_1, o) \text{ in } t \rangle \rightsquigarrow n_1\langle v \rangle \parallel n_2\langle \text{let } x : T = v \text{ in } t \rangle$	CLAIM _i ³	
$\frac{t_2 \neq v}{n_2\langle t_2 \rangle \parallel n_1\langle \text{let } x : T = \text{claim}@ (n_2, o) \text{ in } t'_1 \rangle \rightsquigarrow}$	CLAIM _i ⁴	
$n_2\langle t_2 \rangle \parallel n_1\langle \text{let } x : T = \text{release}(o); \text{get}@n_2 \text{ in } \text{grab}(o); t'_1 \rangle$		
$n_1\langle v \rangle \parallel n_2\langle \text{let } x : T = \text{get}@n_1 \text{ in } t \rangle \rightsquigarrow n_1\langle v \rangle \parallel n_2\langle \text{let } x : T = v \text{ in } t \rangle$	GET _i ²	
$n'\langle \text{set } v \mapsto n_1 \rangle \parallel n_2\langle \text{let } x : T = \text{get}@n_1 \text{ in } t \rangle \rightsquigarrow n_1\langle \text{set } v \mapsto n_1 \rangle \parallel n_2\langle \text{let } x : T = v \text{ in } t \rangle$	GET _i	
$n\langle \text{suspend}(o); t \rangle \rightsquigarrow n\langle \text{release}(o); \text{grab}(o); t \rangle$	SUSPEND	
$o[c, F, \perp] \parallel n\langle \text{grab}(o); t \rangle \xrightarrow{\tau} o[c, F, \top] \parallel n\langle t \rangle$	GRAB	
$o[c, F, \top] \parallel n\langle \text{release}(o); t \rangle \xrightarrow{\tau} o[c, F, \perp] \parallel n\langle t \rangle$	RELEASE	

Table 5: Internal steps

$$\begin{array}{l}
\mathbf{0} \parallel C \equiv C \quad C_1 \parallel C_2 \equiv C_2 \parallel C_1 \quad (C_1 \parallel C_2) \parallel C_3 \equiv C_1 \parallel (C_2 \parallel C_3) \\
C_1 \parallel \nu(n:T).C_2 \equiv \nu(n:T).(C_1 \parallel C_2) \quad \nu(n_1:T_1).\nu(n_2:T_2).C \equiv \nu(n_2:T_2).\nu(n_1:T_1).C
\end{array}$$

Table 6: Structural congruence

Next we show that the type system indeed assures what it is supposed to, most importantly that a promise is indeed fulfilled only once. First we characterize as erroneous situations where a promise is about to be written a second time: A configuration C contains a *write error* if it is of the form $C \equiv \nu(\Theta').(C' \parallel n'\langle \text{let } x : T = \text{bind } t_1 : T_1 \hookrightarrow n \text{ in } t_2 \rangle \parallel n(t))$. Configurations without such write-errors are called *write-error free*, denoted $\vdash C : \text{ok}$. In [53], an analogous condition is called *handle error*.

First we show that a well-typed component does not contain a manifest write-error.

Lemma 2.2. *If $\Delta \vdash C : \Theta$, then $\vdash C : \text{ok}$.*

Proof. By induction on the typing derivations for judgments on the level of components, i.e., for judgments of the form $\Delta \vdash C : \Theta$; the subordinate typing rules from Tables 3 and 4 on the level of threads and expressions do not play a role for the proof. The empty component in the base case of T-EMPTY is clearly write-error free. The cases for the T-NU-rules by straightforward induction. The case T-SUB for subsumption is likewise follows by induction. The cases for T-NCLASS, T-NOBJ, and T-NFUTURE are trivially satisfied, as they mention a single, basic component, only.

Case: T-PAR

We are given $\Delta, \Theta_2 \vdash C_1 : \Theta_1$ and $\Delta, \Theta_1 \vdash C_2 : \Theta_2$. By induction, both C_1 and C_2 are write-error free. The non-trivial case (which we will lead to a contradiction) is when one of the components attempts to write to a promises and the partner already has fulfilled it. So, wlog. assume that $C_1 = \nu(\Theta'_1).(C'_1 \parallel n_1\langle \text{let } x : T = \text{bind } x : T \hookrightarrow n_2 \text{ in } t'' \rangle)$ and $C_2 = \nu(\Theta'_2).(C'_2 \parallel n_2\langle t_2 \rangle)$. Assume that n_2 neither occurs in Θ'_1 nor in Θ'_2 , otherwise no write error is present (since in that case, the name n_2 mentioned on both sides of the parallel refer to different entities). For C_1 to be well-typed, we have $\Delta, \Theta_2 \vdash n : [T_1]^{+-}$ for some type T_1 . For C_2 to be well-typed, we have $\Theta_2 \vdash n : [T_2]^+$ for some type T_2 . Thus, $\Delta \vdash C_1 \parallel C_2 : \Theta_1, \Theta_2$ cannot be derived, which contradicts the assumption. \square

The next standard property shows preservation of well-typedness under internal

$$\begin{array}{ccc}
\frac{C \equiv \rightsquigarrow \equiv C'}{C \rightsquigarrow C'} & \frac{C \rightsquigarrow C'}{C \parallel C'' \rightsquigarrow C' \parallel C''} & \frac{C \rightsquigarrow C'}{\nu(n:T).C \rightsquigarrow \nu(n:T).C'} \\
\frac{C \equiv \overset{\tau}{\rightarrow} \equiv C'}{C \overset{\tau}{\rightarrow} C'} & \frac{C \overset{\tau}{\rightarrow} C'}{C \parallel C'' \overset{\tau}{\rightarrow} C' \parallel C''} & \frac{C \overset{\tau}{\rightarrow} C'}{\nu(n:T).C \overset{\tau}{\rightarrow} \nu(n:T).C'}
\end{array}$$

Table 7: Reduction modulo congruence

reduction. The necessary ancillary lemmas will in general proceed by induction on the typing derivations for judgments of the form $\Delta \vdash C : \Theta$. From a proof-theoretical (and algorithmic) point of view, the type system as formalized in Tables 2, 3, and 4 has an unwelcome property: it is too “non-deterministic” in that it allows the non-structural subsumption rules T-SUB on the level of threads t and on the level of components C at any point in the derivation. This liberality is unwelcome for proofs by induction on the typing derivation as one loses knowledge about the structure of the premises of an applied rule in the derivation.

Lemma 2.3 (Minimal typing). *1. If $\Delta \vdash_m C : \Theta$ and $\Delta' \vdash C : \Theta'$, then $\Delta \leq \Delta'$ and $\Theta \leq \Theta'$.*

2. If $\Delta \vdash_m C : \Theta$ then $\Delta \vdash C : \Theta$.

3. If $\Delta' \vdash C : \Theta'$, then $\Delta \vdash_m C : \Theta$ with $\Delta \vdash \Delta'$ and $\Theta \leq \Theta'$.

Proof. Straightforward. □

Lemma 2.4 (Subject reduction: \equiv). *If $\Delta \vdash_m C_1 : \Theta$ and $C_1 \equiv C_2$, then $\Delta \vdash_m C_2 : \Theta$.*

Proof. We show preservation of typing by the axioms of Table 6. Proceed by induction on the derivation of $\Delta \vdash_m C_1 : \Theta$.

Case: $C \parallel \mathbf{0} \equiv C$ (idempotence)

We are given $\Delta \vdash C \parallel \mathbf{0} : \Theta$. Inverting T-PAR and by T-EMPTY we get as sub-goals $\Delta, \Theta \vdash_m \mathbf{0} : ()$ and $\Delta \vdash_m C : \Theta$, which concludes the case.

Case: $C \equiv C \parallel \mathbf{0}$ (idempotence)

Immediate using T-PAR and T-EMPTY.

Case: $C_1 \parallel C_2 \equiv C_2 \parallel C_1$ (commutativity)

Immediate.

Case: $C_1 \parallel (C_2 \parallel C_3) \equiv (C_1 \parallel C_2) \parallel C_3$ and vice versa (associativity)

By straightforward induction.

Case: $C_1 \parallel \nu(n:T).C_2 \equiv \nu(n:T).(C_1 \parallel C_2)$

where n does not occur free in C_1 . We are given $\Delta \vdash C_1 \parallel \nu(n:T).C_2 : \Theta_1, \Theta_2$, where n neither occurs in Θ_1 nor Θ_2 . Inverting T-PAR and T-NU_f or T-NU_o², we obtain as two subgoals $\Delta, \Theta_2 \vdash C_1 : \Theta_1$ and $\Delta, \Theta_1 \vdash C_2 : \Theta_1, \Theta_2, n:T$, and the result follows by T-PAR and the respective T-NU-rule. The case for T-NU_o¹ works analogously.

Case: $\nu(n_1:T_1).\nu(n_2:T_2).C \equiv \nu(n_2:T_2).\nu(n_1:T_1).C$

Analogously. □

□

Lemma 2.5 (Subject reduction: $\xrightarrow{\tau}$ and \rightsquigarrow). *Assume $\Delta \vdash C : \Theta$.*

1. If $C \xrightarrow{\tau} \acute{C}$, then $\Delta \vdash \acute{C} : \Theta$.

2. If $C \rightsquigarrow \acute{C}$, then $\Delta \vdash \acute{C} : \Theta$.

Proof. The reduction rules of Table 5 are all of the form $C_1 \parallel n\langle t_1 \rangle \xrightarrow{\tau} C_2 \parallel n\langle t_2 \rangle$, where often $C_1 = C_2$ or C_1 and C_2 missing. In the latter case, it suffices to show that $;\Delta, n:[T]^+ \vdash_m t_1 : T$ implies $;\Delta, n:[T]^+ \vdash_m t_2 : T$.

Case: RED: $\text{let } x : T = v \text{ in } t \rightsquigarrow t[v/x]$

By preservation of typing under substitution.

The 5 rules for let and for conditionals are straightforward. The case for stop follows from the fact that stop has every type (cf. rule T-STOP).

Case: PROM: $n' \langle \text{let } x : T' = \text{promise } T \text{ in } t \rangle \rightsquigarrow v(n : T').(n' \langle \text{let } x : T' = n \text{ in } t \rangle)$

The type system assures that $T' = [T]^{+-}$, i.e., for the left-hand side we obtain as subgoal (inverting T-NTHREAD, T-LET, and T-PROM) $x : [T]^{+-}; \Delta, n' : [T]^{+-} \vdash t : T$. The result follows from T-NU, T-LET, and T-NTHREAD.

Case: BIND_i: $n_1 \langle t \rangle = n_1 \langle \text{let } x : T = \text{bind } o.l(\vec{v}) : T_2 \hookrightarrow n_2 \text{ in } t_1 \rangle \xrightarrow{\tau} n_1 \langle \text{let } x : T = () \text{ in } t_1 \rangle \parallel v(n' : \text{Unit}).(n' \langle \text{let } x : T_2 = M.l(o)(\vec{v}) \text{ in } \text{set } x \mapsto n_2 \rangle)$

The type system assures (cf. T-BIND) that $T = [T']^+$ for some type T' . By assumption, we are given $\Delta \vdash n_1 \langle t \rangle : \Theta$ which implies $\Theta = n_1 : [T_1]^+$. Inverting rule T-NTHREAD gives

$$\frac{\frac{\frac{\Delta', n_2 : [T_2]^+ \vdash M.l(o)(\vec{v}) : T_1}{\Delta', n_2 : [T_2]^+ \vdash \text{bind } o.l(\vec{v}) : T_2 \hookrightarrow n_2 : T} \text{T-BIND} \quad x : T \Delta', n_2 : [T_2]^+ \vdash t_1 : T_1 \quad x : T; \hat{\Delta}, n_2 : [T_2]^+}{\Delta', n_2 : [T_2]^+ \vdash \text{let } x : T = \text{bind } o.l(\vec{v}) : T_2 \hookrightarrow n_2 \text{ in } t_1 : T_1} \text{T-LET}}{\Delta \vdash n_1 \langle \text{let } x : T = \text{bind } o.l(\vec{v}) : T_2 \hookrightarrow n_2 \text{ in } t_1 \rangle : n_1 : [T_1]^+, n_2 : [T_2]^+}$$

Rule T-BIND implies that $\Delta = \Delta', n_2 : [T]^{+-}$, i.e., the thread has write permission on n_2 in the pre-state. Furthermore, $\hat{\Delta} \vdash n_2 : [T]^+$, i.e., in the post-state, the thread has lost its write-permission (as it has executed it). In addition, $\hat{\Gamma}$ is empty. With T-PAR we obtain the following two sub-goals.

$$\frac{\Delta', n_2 : [T_2]^{+-} \vdash n_1 \langle \text{let } x : \text{Unit} = () \text{ in } t_1 \rangle : n_1 : [T_1]^+ \quad \Delta', n_1 : [T_1]^+ \vdash v(n' : \text{Unit}).(n' \langle \text{let } x : T_2 = M.l(o)(\vec{v}) \text{ in } \text{set } x \mapsto n_2 \rangle) : n_2 : [T_2]^{+-}}{\Delta \vdash n_1 \langle \text{let } x : \text{Unit} = () \text{ in } t_1 \rangle \parallel v(n' : \text{Unit}).(n' \langle \text{let } x : T_2 = M.l(o)(\vec{v}) \text{ in } \text{set } x \mapsto n_2 \rangle) : n_1 : [T_1]^+, n_2 : [T_2]^{+-}}$$

Both can be straightforwardly solved using T-NFUTURE, T-NU, T-UNIT, T-LET, T-COMPL, and the assumptions.

The remaining rules work similarly. □ □

Lemma 2.6 (Subject reduction: \equiv). *If $\Delta \vdash C_1 : \Theta$ and $C_1 \equiv C_2$, then $\Delta \vdash C_2 : \Theta$.*

Proof. Assume $\Delta \vdash C_1 : \Theta$ and $C_1 \equiv C_2$. By Lemma 2.3(3), $\Delta' \vdash_m C_1 : \Theta'$ s.t. $\Delta \leq \Delta'$ and $\Theta' \leq \Theta$. By Lemma 2.4, $\Delta' \vdash_m C_2 : \Theta'$, and hence by Lemma 2.3(2), also $\Delta' \vdash C_2 : \Theta'$, and the result follows by subsumption (rule T-SUB). □ □

Lemma 2.7 (Subject reduction: $\xrightarrow{\tau}$ and \rightsquigarrow). *Assume $\Delta \vdash C : \Theta$.*

1. *If $C \xrightarrow{\tau} \hat{C}$, then $\Delta \vdash_m \hat{C} : \Theta$.*
2. *If $C \rightsquigarrow \hat{C}$, then $\Delta \vdash_m \hat{C} : \Theta$.*

Proof. As consequence of the corresponding property for minimal typing from Lemma 2.5 and Lemma 2.3. □ □

Lemma 2.8 (Subject reduction). *If $\Xi \vdash C$ and $C \Longrightarrow \hat{C}$, then $\Xi \vdash \hat{C}$.*

Proof. A consequence of Lemma 2.6 and 2.7. □ □

A direct consequence is that all reachable configurations are write-error free:

Corollary 2.9. *If $\Delta \vdash C : \Theta$ and $C \Longrightarrow \hat{C}$, then $\vdash \hat{C} : \text{ok}$.*

Proof. A consequence of Lemma 2.2 and subject reduction from Lemma 2.8. □ □

$\gamma ::= n\langle call\ o.l(\vec{v}) \rangle \mid n\langle get(v) \rangle \mid \nu(n:T).\gamma$	basic labels
$a ::= \gamma? \mid \gamma!$	receive and send labels

Table 8: Labels

2.3.2 External semantics

The external semantics formalizes the environment interaction of an open component as labeled transitions between judgments of the form

$$\Delta \vdash C : \Theta, \quad (3)$$

where Δ represent the *assumptions* about the environment of the component C and Θ the *commitments*. The assumptions require the existence of *named entities* in the environment (plus giving static typing information). The semantics maintains as invariant that the assumption and commitment contexts are disjoint concerning the names for objects, classes, and threads. In addition, the interface keeps information about whether the value of the future n is already known at the interface. If it is, we write $n:T = v$ as binding of the context. We write furthermore $\Delta \vdash n = v$, if Δ contains the corresponding value information and write $\Delta \vdash n = \perp$, if that is not the case. This extension makes the value of a future (once claimed) available at the interface. With these judgments, the external transitions are of the form:

$$\Delta \vdash C : \Theta \xrightarrow{a} \hat{\Delta} \vdash \hat{C} : \hat{\Theta}. \quad (4)$$

Notation 2.10. We abbreviate the tuple of name contexts Δ, Θ as Ξ . Furthermore we understand $\hat{\Delta}, \hat{\Theta}$ as $\hat{\Xi}$, etc.

The labels of the external transitions represent the corresponding interface interaction (cf. Table 8). A component exchanges information with the environment via *call* labels γ_c and *get* labels γ_g . Interaction is either incoming or outgoing, indicated by $?$, resp., $!$. In the labels, n is the identifier of the thread carrying out the call resp. of being queried via *claim* or *get*. Besides that, object and future names (but no class names) may appear as arguments in the communication. Scope extrusion of names across the interface is indicated by the ν -binder. Given a basic label $\gamma = \nu(\Xi).\gamma'$ where Ξ is a name context such that $\nu(\Xi)$ abbreviates a sequence of single $n:T$ bindings (whose names are assumed all disjoint, as usual) and where γ' does not contain any binders, we call γ' the *core* of the label and refer to it by $\lfloor \gamma \rfloor$. We define core analogously for receive and send labels. The free names $fn(a)$ and the bound names $bn(a)$ of a label a are as usual, whereas $names(a)$ refer to all names of a . In addition, we distinguish between names occurring as arguments of a label, in *passive* position, and the name occurring as carrier of the activity, in *active* position. Name n , for illustration, occurs actively and free in $n\langle call\ o.l(\vec{v}) \rangle$ and in $n\langle get(v) \rangle$. We write $fn_a(a)$ for the free names occurring in active position, $fn_p(a)$ for the free names in passive position, etc. All notations are used analogously for basic labels γ . Note that for incoming labels, Ξ contains only bindings to environment objects (besides future names), as the environment cannot create component objects; dually for outgoing communication.

The steps of the operational semantics for open systems checks the *static* assumptions, i.e., whether at most the names actually occurring in the core of the label are mentioned in the ν -binders of the label, and whether the transmitted values are of the correct types. This is covered in the following definition.

$\frac{\hat{\Xi} = \hat{\Xi}_1, n:[T]^+, \hat{\Xi}_2 \quad ; \hat{\Xi} \vdash \vec{v} : \vec{T} \quad a = n(\text{call } o, l(\vec{v}))?}{\hat{\Xi} \vdash a : \vec{T} \rightarrow _} \text{LT-CALLI}$	$\frac{; \hat{\Xi} \vdash v : T \quad a = n(\text{get}(v))?}{\hat{\Xi} \vdash a : _ \rightarrow T} \text{LT-GETI}$
---	---

Table 9: Typechecking labels

Definition 2.11 (Well-formedness and well-typedness of a label). A label $a = \nu(\Xi).[a]$ is well-formed, written $\vdash a$, if $\text{dom}(\Xi) \subseteq \text{fn}([a])$ and if Ξ is a well-formed name-context for object and future names, i.e., no name bound in Ξ occurs twice. The assertion

$$\hat{\Xi} \vdash o.l? : \vec{T} \rightarrow T \quad (5)$$

(“an incoming call of the method labeled l in object o expects arguments of type \vec{T} and results in a value of type T ”) is given by the following rule, i.e., implication:

$$\frac{; \hat{\Theta} \vdash o : c \quad ; \hat{\Xi} \vdash c : [(\dots, l:\vec{T} \rightarrow T, \dots)]}{\hat{\Xi} \vdash o.l? : \vec{T} \rightarrow T} \quad (6)$$

For outgoing calls, $\hat{\Xi} \vdash o.l! : \vec{T} \rightarrow T$ is defined dually. In particular, in the first premise, $\hat{\Theta}$ is replaced by $\hat{\Delta}$. Well-typedness of an incoming core label a with expected type \vec{T} , resp., T , and relative to the name context $\hat{\Xi}$ is asserted by

$$\hat{\Xi} \vdash a : \vec{T} \rightarrow _ \quad \text{resp.}, \quad \hat{\Xi} \vdash a : _ \rightarrow T, \quad (7)$$

as given by Table 9. In LT-CALLI, the premise $; \hat{\Xi} \vdash \vec{v} : \vec{T}$ is interpreted in such a way that checking for write-permission consumes that permission (analogous to the corresponding premise of T-BIND in Table 4): Let $\hat{\Xi}_0$ abbreviate $; \hat{\Xi}$. Then $; \hat{\Xi} \vdash \vec{v} : \vec{T}$ means: $\hat{\Xi}_i \vdash v_i : T_i$ and $\hat{\Xi}_{i+1} = \hat{\Xi}_i \setminus T_i$, for all $0 \leq i \leq n-1$.

Note that the receiver o of the call is checked using only the commitment context $\hat{\Theta}$, to assure that o is a component object. Note further that to check the interface type of the class c , the full $\hat{\Xi}$ is consulted, since the argument types \vec{T} or the result type T may refer to both component and environment classes. The incremental type of first premise $; \hat{\Xi} \vdash \vec{v} : \vec{T}$ of LT-CALLI assures that no name is transmitted twice with write-permission. In a similar spirit: requiring that $\hat{\Xi}$ is of the form $\hat{\Xi}_1, n:[T]^+, \hat{\Xi}_2$ assures that it is not possible to transmit n with write-permissions if n is the active thread of the label.

Besides *checking* whether the assumptions are met before a transition, the contexts are *updated* by a transition step, i.e., extended by the new names, whose scope extrudes. For the binding part Ξ' of a label $\nu(\Xi').\gamma$, the scope of the references to existing objects and thread names Δ' extrudes across the border. In the step, Δ' extends the assumption context Δ and Θ' the commitment context Θ . Besides information about new names, the context information is potentially updated wrt. the availability of a future *value*. This is done when a get-label is exchanged at the interface for the first time, i.e., when a future value is claimed successfully for the first time. For outgoing communication, the situation is dual.

Before we come to the corresponding Definition 2.12 below, we make clear (again) the interpretation of judgments $\Delta \vdash C : \Theta$. Interesting is in particular the information $n:[T]^+$, stipulating that name n is available with write-permission (and result type

T). In case of $\Delta \vdash n : [T]^{+-}$, the name n is assumed to be available in the environment as writeable, and conversely $\Theta \vdash n : [T]^{+-}$ asserts write permission for the component. Since read permissions, captured by types $[T]^+$, are not treated linearly—one is allowed to read from a future reference as many times as wished—the treatment of bindings $n:[T]^+$ is simpler. Hence here we concentrate on $n:[T]^{+-}$ and the write permissions.

Since the domains of Δ and Θ are disjoint, bindings $n:T'$ cannot be available in the assumption context Δ and the commitments Θ at the same time. The information $T' = [T]^{+-}$ indicates which side, component or environment, has the write permission. If, for instance $\Delta \vdash n : [T]^{+-}$, then the component is not allowed to execute a bind on reference n . The same restriction does not apply to read permissions. In the mentioned situation $\Delta \vdash n : [T]^{+-}$, the component can execute a *claim*-operation on n . The same applies if $\Delta \vdash n : [T]^+$. In other words, a name n can be accessed by reading by both the environment and the component once known at the interface, independent whether it is part of Δ or of Θ . A difference between bindings of the form $n:[T]^{+-}$ and $n:[T]^+$ (and likewise $n:[T]^+ = v$) is, that communication can *change* $\Delta \vdash n : [T]^{+-}$ to $\Theta \vdash n : [T]^{+-}$ and vice versa. For names n of type $[T]^+$, this change of side is impossible. The latter kind of information, for instance $\Theta \vdash n : [T]^+$, implies that the code has been bound to n and it is placed in the component. Once fixed there, the reference to n may, of course, be passed around, but the thread named n itself cannot change to the environment since the language does not support *mobile* code.

Now, how does communication labels as interface interactions update the binding contexts? We distinguish two ways, the name n of a thread can be transmitted in a label: *passively*, when transported as the argument of a call or a get-interaction, and *actively*, when mentioned as the carrier of the activity, as the n in $n\langle \text{call } o.l(\vec{v}) \rangle$ and $n\langle \text{get}(v) \rangle$. As usual, such references (actively or passively) can be transmitted as fresh names, i.e., under a ν -binder, or alternatively as an already known name. When transmitted *passively* and typed with $[T]^{+-}$ for some type T , the write-permission to n is handed over to the receiving side and at the same time, that permission is removed from the sender side. So if, e.g., the environment is assumed to possess the write-permission for reference n , witnessed by $\Delta \vdash n : [T]^{+-}$, then sending n as argument in a communication to the component removes the binding from the environment and adds the permission to the component side, yielding $\Theta \vdash n : [T]^{+-}$. In case the name is transmitted *actively*, the receiver does not obtain write permission.

Now, what about transmitting n *actively*? An incoming call $n\langle \text{call } o.l(\vec{v}) \rangle$?, e.g., reveals at the interface that the promise indeed has been fulfilled. As, in that situation of an incoming call, the thread is located at the component, the commitment context is updated to satisfy $\Theta \vdash n : [T]^+ = \perp$ (for an appropriate type T) after the communication. Indeed, before the step it is checked, that the environment actually has write permission for n , i.e., that $\Delta \vdash n : [T]^{+-}$, or that the name n is new. See the incoming call in Figure 2(a), where the n is fresh, resp. in 2(c), where the n has been transmitted passively and with write-permissions to the environment before the call (in the dotted arrow).

Whereas call-labels make public, at which side the thread in question resides, get-labels, on the other hand, reveal that the thread has terminated and fix the result value (if that information had not been public interface information before). There are two situations, where a, say, outgoing get-communication is possible. In both cases, the future resides in the component and after the get-communication, the value is determined, i.e., $\Theta \vdash n : [T]^+ = v$ (if not already before the step). One scenario is that $\Delta \vdash n : [T]^+ = \perp$ before the step still. If, in that situation, the *get* is executed by

the environment, it is required that the component must have had write permission before that step, i.e., $\Theta \vdash n : [T]^{+-}$ (cf. Figure 2(b)). The only way, the value for n is available for the environment now is that, *unnoticed*⁶ at the interface, the promise had been fulfilled and the corresponding thread already has terminated, and this could have been done by the component, only. In that situation, the contexts are updated from $\Theta \vdash n : [T]^{+-}$ to $\Theta \vdash n : [T]^+ = v$: the component loses the write-permission as it obviously has executed its permission already and the value v is fixed and known at the interface. Alternatively, the thread may be known to be part of the component with the promise already fulfilled ($\Theta \vdash n : [T]^+ = \perp$, as shown in Figure 2(a) and 2(c)). Finally, the value for n might already be known at the interface, i.e., already before the step, $\Theta \vdash n : [T]^+ = v$ holds. In that situation, v has been added as interface information previously, either by a prior get-interaction incoming get-communication or an outgoing return-communication, and the situation corresponds to the last get in Figure 2(b) and 2(c).




Figure 2: Scenarios

This gives rise to the following definition.

Definition 2.12 (Context update). Let Ξ be a name context and $a = v(\Xi').[a]$ an incoming label. We define the (intermediate) contexts $\Theta'' = \Theta$ and $\Delta'' = \Delta, \Xi'$.

Let furthermore Σ'' be the set of bindings defined as follows. In case of a call label, i.e., $[a] = n\langle \text{call } o.l(\vec{v}) \rangle?$, let the vector of types \vec{T} be defined by $\Xi \vdash o.l? : \vec{T} \rightarrow T$ according to equation (5) of Definition 2.11. Then Σ'' consists of bindings of the form $v_i : [T_i]^{+-}$ for values v_i from \vec{v} such that $T_i = [T_i]^{+-}$. In case of a get label, i.e., $[a] = n\langle \text{get}(v) \rangle?$, the context Σ'' is $v : [T]^+ = v$ if $\Delta'' \vdash n : [[T]^+]$, and empty otherwise.

With Σ'' given this way, the definitions of the post-contexts $\hat{\Delta}$ and $\hat{\Theta}$ distinguish between calls and get-interaction: If a is a call label and $n \in \text{names}_a(a)$, we define

$$\hat{\Delta} = \Delta'' \setminus \Sigma'' \setminus n : [T]^{+-} \quad \text{and} \quad \hat{\Theta} = \Theta'', \Sigma'', n : [T]^+ . \quad (8)$$

If a is a get label $a = v(\Xi').n\langle \text{get}(v) \rangle?$ and $n \in \text{names}_a(a)$, $\hat{\Delta}$ and $\hat{\Theta}$ are given by:

$$\hat{\Delta} = \Delta'' \setminus \Sigma'', n : [T]^+ = v \quad \text{and} \quad \hat{\Theta} = \Theta'', \Sigma'' . \quad (9)$$

For outgoing communication, the definition is applied dually.

The definition proceeds in two steps. In a first step, the assumption and the commitment contexts Δ and Θ are extended with the bindings Ξ' carried with the incoming

⁶It is important that the bind-operation on a promise is an internal action and *not* recorded at the interface. This is also the reason to represent an evaluated future n by $n'\langle \text{set } v \mapsto n \rangle$, where n' is hidden behind a v -binder and not by $n\langle v \rangle$ (cf. rule BIND_i of Table 5).

label a . Note that the bindings $\Xi' \vdash n : [T]^{+-}$ or $\Xi' \vdash n : [T]^+$ for future references, kept in Σ' , are added to the assumption context Δ but not the commitment context (in the considered case of incoming communication). The second step deals with the write permissions, i.e., it transfers the write permission transmitted arguments from the sender side to the receiver side. The binding context Σ'' deals with the permissions carried by thread names transmitted passively, i.e., as arguments of the communication. It remains to take care also of the information carried by the active thread. For that we distinguish calls and get-labels. An incoming call (equation (8)) with n as active thread additionally is the sign that the thread is now located at the component side and that the write permission has been consumed by the environment side. Hence, in equation (8), the environment loses the write-permission and the component is extended by the binding $n:[T]^+$. In case of an incoming get, the transmitted value v is remembered as part of Δ (cf. equation (8)).

Now to the interface behavior. Corresponding to the labels from Table 8, there are a number of rules for external communication: either incoming or outgoing calls, resp., get-labels. All rules have some premises in common. In all cases, the context Ξ before the interaction is updated to $\hat{\Xi} = \Xi + a$ using Definition 2.12, where a is the interaction label. The rules for incoming communication differ from the corresponding ones for outgoing communication in that well-typedness and well-formedness of the label is checked by the premises $\hat{\Xi} \vdash [a] : \vec{T} \rightarrow _$, resp. $\hat{\Xi} \vdash [a] : _ \rightarrow \vec{T}$ (for calls) resp., $\hat{\Xi} \vdash [a] : _ \rightarrow T$ (for get-labels), using Definition 2.11. For outgoing communication, the check is unnecessary as starting with a well-typed component, there is no need in re-checking now, as the operational steps preserve well-typedness (subject reduction).

When the component claims the value of a future, we distinguish two situations: the future value is accessed for the first time across the interface or not. In the first case, corresponding to rules CLAIM₁ and CLAIM₂, the interface does not contain the value of the future yet, stipulated by the premise $\Delta \vdash n' = \perp$. Remember that $\Delta \vdash n$ requires that the thread n is part of the environment. In that situation it is unclear from the perspective of the component, whether or not the value has already been computed. Hence, it is possible that executing *claim* is immediately successful (cf. rule CLAIM₁) or that the thread n trying to obtain the value has to suspend itself and try later (cf. rule CLAIM₂). The external rule CLAIM₂ works exactly like the corresponding internal rule CLAIM_i² from Table 5, except that here it is required that the queried future n' is part of the environment. The behavior of a thread wrt. claiming a future value is illustrated in Figure 1. If the future value is already known at the interface (cf. rule CLAIM₃ and especially premise $\Delta \vdash n' = v$), executing *claim* is always successful and the value v is (re-)transmitted. *get* works analogously to *claim*, except that *get* insists of obtaining the value, i.e., the alternative of relinquishing the lock and trying again as in rule CLAIM₂, is not available for *get*. The last two rules deal with the situation that the environment fetches the value.

Finally, we characterize the *initial* situation. Initially, the component contains at most one initial activity and no objects. More precisely, given that $\Xi_0 \vdash C_0$ is the initial judgment, then C_0 contains no objects. Concerning the threads as the active entities: initially exactly one thread is executing, either at the component side or at the environment side. The distinction is made at the interface that initially either $\Theta_0 \vdash n$ or $\Delta_0 \vdash n$, where n is the only thread name in the system.

Remark 2.13 (Comparison with Java-like multi-threading). *The formalization for the multi-threaded case, for instance in [4], is quite similar. One complication encountered there is that one has to take reentrance into account. The rule for incoming call*

$\frac{a = v(\Xi'), n(\text{call } o.l(\vec{v}))? \quad \hat{\Xi} = \Xi + a \quad (\Xi' \vdash n \vee \Delta \vdash n : [_]^{+-}) \quad \hat{\Xi} \vdash o.l? : \vec{T} \rightarrow T \quad \hat{\Xi} \vdash [a] : \vec{T} \rightarrow _}{\Xi \vdash C \xrightarrow{a} \hat{\Xi} \vdash C \parallel n(\text{let } x:T = \text{grab}(o); M.l(o)(\vec{v}) \text{ in } \text{release}(o); x)}$	CALLI
$\frac{a = v(\Xi'), n(\text{call } o.l(\vec{v}))! \quad \Xi' = \text{fn}([a]) \cap \Xi_1 \quad \hat{\Xi}_1 = \Xi_1 \setminus \Xi' \quad \Delta \vdash o \quad \hat{\Xi} = \Xi + a}{\Xi \vdash v(\Xi_1).(C \parallel n(\text{let } x:T = \text{bind } o.l(\vec{v}) : T \hookrightarrow n \text{ in } t)) \xrightarrow{a} \hat{\Xi} \vdash v(\hat{\Xi}_1).(C \parallel n(\text{let } x : T = n \text{ in } t))}$	CALLO
$\frac{a = v(\Xi'), n(\text{get}(v))? \quad \hat{\Xi} = \Xi + a \quad \Delta \vdash n' = \perp \quad \hat{\Xi} \vdash [a] : _ \rightarrow T}{\Xi \vdash v(\Xi_1).(C \parallel n(\text{let } x:T = \text{claim}@(\underline{n'}, _) \text{ in } t)) \xrightarrow{a} \hat{\Xi} \vdash v(\Xi_1).(C \parallel n(\text{let } x:T = v \text{ in } t))}$	CLAIMI ₁
$\frac{\Delta \vdash n' = \perp}{\Xi \vdash v(\Xi_1).(C \parallel n(\text{let } x:T = \text{claim}@(\underline{n'}, o) \text{ in } t)) \rightsquigarrow \Xi \vdash v(\Xi_1).(C \parallel n(\text{let } x : T = \text{release}(o); \text{get}@n' \text{ in } \text{grab}(o); t))}$	CLAIMI ₂
$\frac{a = n'(\text{get}(v))? \quad \Delta \vdash n' = v \quad \Xi \vdash [a] : _ \rightarrow T}{\Xi \vdash v(\Xi_1).(C \parallel n(\text{let } x:T = \text{claim}@(\underline{n'}, _) \text{ in } t)) \xrightarrow{a} \Xi \vdash v(\Xi_1).(C \parallel n(\text{let } x:T = v \text{ in } t))}$	CLAIMI ₃
$\frac{a = v(\Xi'), n'(\text{get}(v))? \quad \hat{\Xi} = \Xi + a \quad \Delta \vdash n' = \perp \quad \hat{\Xi} \vdash [a] : _ \rightarrow T}{\Xi \vdash v(\Xi_1).(C \parallel n(\text{let } x:T = \text{get}@n' \text{ in } t)) \xrightarrow{a} \hat{\Xi} \vdash v(\Xi_1).(C \parallel n(\text{let } x:T = v \text{ in } t))}$	GETI ₁
$\frac{a = n'(\text{get}(v))? \quad \Delta \vdash n' = v \quad \Xi \vdash [a] : _ \rightarrow T}{\Xi \vdash v(\Xi_1).(C \parallel n(\text{let } x:T = \text{get}@n' \text{ in } t)) \xrightarrow{a} \Xi \vdash v(\Xi_1).(C \parallel n(\text{let } x:T = v \text{ in } t))}$	GETI ₂
$\frac{a = v(\Xi'), n(\text{get}(v))! \quad \Xi' = \text{fn}([a]) \cap \Xi_1 \quad \hat{\Xi}_1 = \Xi_1 \setminus \Xi' \quad \hat{\Xi} = \Xi + a}{\Xi \vdash v(\Xi_1).(C \parallel v(n':T).(n'(\text{set } v \mapsto n))) \xrightarrow{a} \hat{\Xi} \vdash v(\hat{\Xi}_1).(C \parallel n(v))}$	GETO ₁
$\frac{a = n(\text{get}(v))! \quad \Theta \vdash n = v}{\Xi \vdash C \xrightarrow{a} \Xi \vdash C}$	GETO ₂

Table 10: External steps

CALLI in Table 10 deals with a non-reentrance situation, which is the only situation relevant in the setting here. In addition to the rule CALLI, Java-like multi-threading requires further CALLI-rules to cover the situations, when the call is reentrant. \square

3 Interface behavior

Next we characterize the possible (“legal”) *interface behavior* as interaction traces between component and environment. Half of the work has been done already in the definition of the external steps in Table 10: For incoming communication, for which the environment is responsible, the assumption contexts are consulted to check whether the communication originates from a realizable environment. Concerning the reaction of the component, no such checks were necessary. To characterize when a given trace is *legal*, the behavior of the component side, i.e., the outgoing communication, must adhere to the dual discipline we imposed on the environment for the open semantics. This means, we analogously abstract away from the program code, rendering the situation symmetric.

$\Xi \vdash \epsilon : \text{trace}$	L-EMPTY
$\frac{a = v(\Xi'), n(\text{call } o.l(\vec{v}))? \quad \hat{\Xi} = \Xi + a \quad (\Xi' \vdash n \vee \Delta \vdash n : \mathbb{I}^{++})}{\hat{\Xi} \vdash o.l? : \vec{T} \rightarrow T \quad \hat{\Xi} \vdash [a] : \vec{T} \rightarrow _ \quad \hat{\Xi} \vdash s : \text{trace}} \text{L-CALLI}$	
$\Xi \vdash a s : \text{trace}$	
$\frac{a = v(\Xi'), n(\text{get}(v))? \quad \hat{\Xi} = \Xi + a \quad \Delta \vdash n = \perp \quad \hat{\Xi} \vdash [a] : _ \rightarrow T \quad \hat{\Xi} \vdash s : \text{trace}}{\Xi \vdash a s : \text{trace}} \text{L-GETI}_1$	
$\Xi \vdash a s : \text{trace}$	
$\frac{a = n(\text{get}(v))? \quad \Delta \vdash n = v \quad \Xi \vdash s : \text{trace}}{\Xi \vdash a s : \text{trace}} \text{L-GETI}_2$	
$\Xi \vdash a s : \text{trace}$	

Table 11: Legal traces (dual rules omitted)

3.1 Legal traces system

The rules of Table 11 specify legality of traces. We use the same conventions and notations as for the operational semantics (cf. Notation 2.10). The judgments in the derivation system are of the form

$$\Xi \vdash s : \text{trace} . \quad (10)$$

We write $\Xi \vdash t : \text{trace}$, if there exists a derivation according to the rules of Table 11 with an instance of L-EMPTY as axiom. The empty trace is always legal (cf. rule L-EMPTY), and distinguishing according to the first action a of the trace, the rules check whether a is possible. Furthermore, the contexts are updated appropriately, and the rules recur checking the tail of the trace. The rules are symmetric wrt. incoming and outgoing communication (the dual rules are omitted). Rule L-CALLI for incoming calls works completely analogously to the CALLI-rule in the semantics: the second premise updates the context Ξ appropriately with the information contained in a , premise $\Xi' \vdash n$ of L-CALLI assures that the identity n of the future, carrying out the call, is fresh and the two premises $\hat{\Xi} \vdash o.l? : \vec{T} \rightarrow _$ and $\hat{\Xi} \vdash [a] : \vec{T} \rightarrow _$ together assure that the transmitted values are well-typed (cf. Definition 2.11); the latter two checks correspond to the analogous premises for the external semantics in rule CALLI, except that the return type of the method does not play a role here. The L-GETI-rules for claiming a value work similarly. In particular the type checking of the transmitted value is done by the combination of the premises $\Delta \vdash n : [T]$ and $\hat{\Xi} \vdash [a] : _ \rightarrow T$. As in the external semantics, we distinguish two cases, namely whether the value of the future has been incorporated in the interface already or not (rules L-GETI₂ and L-GETI₁). In both cases, the thread must be executing on the side of the environment for an incoming get. This is checked by the premise $\Delta \vdash n = \perp$ resp. by $\Delta \vdash n = v$. In case of L-GETI₂, where the value of the future has been incorporated as v into the interface information, the actual parameter of the get-label must, of course, be v . If not (for L-GETI₁), the transmitted argument value is arbitrary, apart from the fact that it must be consistent with the static typing requirements.

It remains to show that the behavioral description, as given by Table 11, actually does what it claims to do, to characterize the possible interface behavior of well-typed components. We show the soundness of this abstraction plus the necessary ancillary lemmas such as subject reduction.

Lemma 3.1 (Subject reduction). $\Xi_0 \vdash C \xrightarrow{s} \hat{\Xi} \vdash \hat{C}$, then $\hat{\Xi} \vdash \hat{C}$.

Proof. By induction on the number of reduction steps. That internal steps preserve well-typedness, i.e., $\Xi \vdash C \implies \Xi \vdash C$, follows from Lemma 2.8. That leaves the external reduction steps of Table 10.

Case: CALLI

We are given $\Xi \vdash C$. The disjunctive premise of the rule distinguishes two sub-cases: 1) $\Xi' \vdash n$ (the thread name is transmitted freshly) or 2) $\Delta \vdash n : [_]^{+-}$ (the thread is not transmitted freshly and the environment has write-permission before the step). Both are treated uniformly in the argument. For the right-hand side of the transition, we need to show

$$\hat{\Xi} \vdash C \parallel n \langle \text{let } x:T = \text{grab}(o); M.l(o)(\vec{v}) \text{ in } \text{release}(o); x \rangle .$$

According to the definition of context update (Definition 2.12), $\hat{\Xi} = \hat{\Delta}, \hat{\Theta}$, where $\hat{\Theta} = \Theta, \Sigma'', n : [T]^+$ and where Σ'' contains bindings $n':[T']^{+-}$ for those references transmitted with read-write permission as argument of the call (see the right-hand of equation (8)). The assumption context $\hat{\Delta}$ for \hat{C} after the step (by the left-hand of the same equation) is of the form $\Delta, \Delta', \Sigma' \setminus \Sigma'' \setminus n:[T]^{+-}$, which we abbreviate by $\Delta, \Delta', \Sigma_\Delta$. So for the new thread n at component side, we need to show that

$$\Delta, \Delta', \Sigma_\Delta, \Theta \vdash n \langle \text{let } x:T = \text{grab}(o); M.l(o)(\vec{v}) \text{ in } \text{release}(o); x \rangle : n:[T]^+, \Sigma'' . \quad (11)$$

This follows by rules T-NFUTURE, T-LET, T-GRAB, preservation of typing under substitution, T-RELEASE, and the axiom T-VAR. Note that the result type T (which is the type of x) is guaranteed by the premise $\hat{\Xi} \vdash o.l? : \vec{T} \rightarrow T$ of the reduction rule CALLI. From equation (11), the result follows by T-PAR, subsumption, and the assumption $\Xi \vdash C$.

Case: CALLO

We are given

$$\Xi \vdash v(\Xi_1).(C \parallel n' \langle \text{let } x:T = \text{bind } o.l(\vec{v}) : T \leftrightarrow n \text{ in } t \rangle)$$

before the step. By one of the premises of rule CALLO we know $\Delta \vdash o$, i.e., object o is an environment object ⁷ That o refers to an object is assured by the type system and the assumption that the pre-configuration is well-typed.

We distinguish two sub-cases, namely whether promise n 1) is known at the interface before the step or 2) it is hidden still. In the first case we have $\Theta \vdash n:[T']^{+-}$ with $T = [T']^{+-}$ (as a consequence of the fact that the configuration is well-typed. Especially, inverting T-NFUTURE and T-BIND entails that the component must have write-permission for n to be well-typed). The result follows by the typing rules T-NU, T-PAR, T-LET, and T-NAME.

Case: CLAIM₁

The core of the type preservation here is to assure that the claim-statement in the pre-configuration and the transmitted value v in the post-configuration are of the same appropriate type T . Well-typedness of the pre-configuration implies with $\text{claim}@(\vec{n}', o)$ of type T , that the reference n' is of type $[T]^+$. The third premise of CLAIM₁ states $\hat{\Xi} \vdash [a] : _ \rightarrow T$, which implies with Definition 2.11, especially rule LT-GETI of Table 9, that also v is of type T , as required.

⁷We do not allow cross-border instantiation here, i.e., the component is not allowed to instantiate environment objects and vice versa.

Case: CLAIM₂

By inverting the type rules T-NU, T-PAR, T-LET and T-CLAIM for the pre-configuration of the step, and by using the same typing rules (except T-CLAIM) plus T-GET, T-RELEASE, and T-GRAB.

The remaining rules work similarly. \square \square

Lemma 3.2 (Soundness of abstractions). *If $\Xi_0 \vdash C$ and $\Xi_0 \vdash C \xRightarrow{t}$, then $\Xi_0 \vdash t$: trace.*

Proof. By induction on the number of steps in \xRightarrow{t} . The base case of zero steps (which implies $t = \epsilon$) is immediate, using L-EMPTY. The induction for internal steps of the form $\Xi \vdash C \Longrightarrow \Xi \vdash \dot{C}$ follow by subject reduction for internal steps from Lemma 2.8; in particular, internal steps do not change the context Ξ . Remain the external steps of Table 10. First note the contexts Ξ are *updated* by each external step to $\dot{\Xi}$ the same way as the contexts are updated in the legal trace system.

The cases for incoming communication are checked straightforwardly, as the operational rules check incoming communication for legality, already, i.e., the premises of the operational rules have their counterparts in the rules for legal traces.

Case: CALLI

Immediate, as the premises of L-CALLI coincide with the ones of CALLI.

Case: CLAIM₁ and GET₁

The two cases are covered by rule L-GET₁, which has the same premises as the operational rules.

Case: CLAIM₂

Trivial, as the step is an internal one.

Case: CLAIM₃ and GET₂

The two cases are covered by L-GET₂.

The cases for outgoing communication are slightly more complex, as the label in the operational rule is not type-checked or checked for well-formedness as for incoming communication and as is done in the rules for legality.

Case: CALLO

We need to check whether the premises of L-CALLO, the dual to L-CALLI of Table 11, are satisfied. By assumption, the pre-configuration

$$\Xi \vdash \nu(\Xi_1).(C \parallel n' \langle \text{let } x:T = \text{bind } o.l(\vec{v}) : T \leftrightarrow n \text{ in } t \rangle) \quad (12)$$

is well-typed. For thread name n this implies, it is bound either in Ξ or in Ξ_1 , more precisely, either $\Theta \vdash n : [T]^{+-}$ (it is public interface information that the component has write-permission for n) or $\Xi_1 \vdash n : [T]^{+-}$ (the name n is not yet known in the environment before the communication). In the latter situation we obtain $\Xi' \vdash n : []^{+-}$ by the premise $\Xi' = \text{fn}([a]) \cap \Xi_1$ of CALLO. Thus, the third premise $\Xi' \vdash n \vee \Theta \vdash n : []^{+-}$ of L-CALLO is satisfied. We furthermore need to check whether the label is type-correct (checked by premises nr. 4 and 5 or L-CALLO). Its easy to check that the label is well-formed (cf. the first part of Definition 2.11). The first premise of the check of equation (6), that the receiving object o is an environment object, is directly given by the premise $\Delta \vdash o$ of CALLO. That the object o supports a method labeled l (of type $\vec{T} \rightarrow T$) follows from the fact that the pre-configuration of the call-step is well-typed. So this gives L-CALLO's premise $\dot{\Xi} \vdash o.l! : \vec{T} \rightarrow T$. Remains the type check $\dot{\Xi} \vdash [a] : \vec{T} \rightarrow _$ (checking that the transmitted values \vec{T} are of the excepted type \vec{t}), which again follows from well-typedness of equation (12) (especially inverting T-BIND).

The remaining cases work similarly. \square \square

Remark 3.3 (Comparison with reentrant threading). *In a multi-threaded setting with synchronous method calls (see for instance [4] [59]), the definition of legal traces is more complicated. Especially, to judge whether a trace s is possible required referring to the past. I.e., instead of judgments of the form of equation (10), the check for legality with synchronous calls uses judgments of the form:*

$$\Xi \vdash r \triangleright s : \text{trace} ,$$

reading “after history r (and in the context Ξ), the trace s is possible”. This difference has once more to do with reentrance, resp. with the absence of this phenomenon here. In the threaded case, where, e.g., an outgoing call can be followed by a subsequent incoming call as a “call-back”. To check therefore, whether a call or a return is possible as a next step involves checking the proper nesting of the call- and return labels. This nesting requirement (also called the balance condition) degenerates here in the absence of call-backs to the given requirement that each call uses a fresh (future) identity and that each get-label (taking the role of the return label in the multithreaded setting) is preceded by exactly one matching preceding call. This can be judged by $\Delta \vdash n : [-]$ or $\Theta \vdash n : [-]$ (depending on whether we are dealing with incoming or outgoing get-labels) and especially, no reference to the history of interface interactions is needed. \square

Remark 3.4 (Monitors). *The objects of the calculus act as monitors as they allow only one activity at a time inside the object. For the operational semantics of Section 2.3, the lock-taking is part of the internal steps. In other words, the handing-over of the call at the interface and the actual entry into the synchronized method body is non-atomic, and at the interface, objects are input-enabled.*

This formalization therefore resembles the one used for the interface description of Java-like reentrant monitors in [3]. To treat the interface interaction and actual lock-grabbing as non-atomic leads to a clean separation of concerns of the component and of the environment. In [3], this non-atomicity, however, give rise to quite complex conditions characterizing the legal interface behavior. In short, in the setting of [3], it is non-trivial to characterize exactly those situations, when the lock of the object is necessarily taken by one thread which makes certain interactions of other threads impossible. This characterization is non-trivial especially as the interface interaction is non-atomic.

Note, however, that these complications are not present in the current setting with active objects, even if the the objects acts as monitors like in [3]. The reason is simple: there is no need to capture situations when the lock is taken. In Java, the synchronization behavior of a method is part of the interface information. Concretely, the synchronized-modifier of Java, specifies that the method’s body is executed atomically in that object without interference of other⁸ threads, assuming that all other methods of the callee are synchronized, as well. Here, in contrast, there is no interface information that guarantees that a method body is executed atomically. In particular, the method body can give up the lock temporarily via the suspend-statement, but this fact is not reflected in the interface information here. This absence of knowledge simplifies the interface description considerably. \square

⁸Note that a thread can “interfere” in that setting with itself due to recursion and reentrance.

4 Conclusion

We presented an open semantics describing the interface behavior of components in a concurrent object-oriented language with futures and promises. The calculus corresponds to the core of the *Creol* language, including classes, asynchronous method calls, the synchronization mechanism, and futures, and extended by promises. Concentrating on the black-box interface behavior, however, the interface semantics is, to a certain extent, independent of the concrete language and is characteristic for the mentioned features; for instance, extending *Java* with futures (see also the citations below) would lead to a quite similar formalization (of course, low level details may be different). Concentrating on the concurrency model, certain aspects of *Creol* have been omitted here, most notably inheritance and safe asynchronous class upgrades.

Related work

The general concept of “delayed reference” to a result of a computation to be yet completed is quite old. The notion of futures was introduced by Baker and Hewitt [13], where `(future e)` denotes an expression executed in a separate thread, i.e., concurrently with the rest of the program. As the result of the e is not immediately available, a *future variable* (or future) is introduced as placeholder, which will eventually contain the result of e . In the meantime, the future can be passed around, and when it is accessed for reading (“touched” or “claimed”), the execution suspends until the future value is available, namely when e is evaluated. The principle has also been called *wait-by-necessity* [16][17]. Futures provide, at least in a purely functional setting, an elegant means to introduce concurrency and *transparent* synchronization simply by accessing the futures. They have been employed for the parallel *Multilisp* programming language [36].

Indeed, quite a number of calculi and programming languages have been equipped with concurrency using future-like mechanisms and asynchronous method calls. Flanagan and Felleisen [31] [29] [30] present an operational semantics (based on evaluation contexts) for a λ -calculus with futures. The formalization is used for an analysis and optimization technique to eliminate superfluous dereferencing (“touches”) of future variables. The analysis is an application of a set-based analysis and the resulting transformation is known as touch optimization. Moreau [51] presents a semantics of Scheme equipped with futures and control operators. *Promises* is a mechanism quite similar to futures and actually the two notions are sometimes used synonymously. They have been proposed in [48]. A language featuring both futures and promises as separate concepts, is *Alice ML* [9][45][58].

[53] presents a concurrent call-by-value λ -calculus with reference cells (i.e., a non-purely functional calculus with an imperative part and a heap) and with futures (λ_{fut}), which serves as the core of *Alice ML* [9] [57] [45]. Certain aspects of that work are quite close to the material presented here. In particular, we were inspired by using a type system to avoid fulfilling a promise twice (in [53] called handle error). There are some notable differences, as well. The calculus incorporates futures and promises into a λ -calculus, such that functions can be executed in parallel. In contrast, the notion of futures here, in an object-oriented setting, is coupled to the asynchronous execution of methods. Furthermore, the object-oriented setting here, inspired by *Creol*, is more high-level. In contrast, λ_{fut} relies on an atomic test-and-set operation when accessing the heap to avoid atomicity problems. Besides that, they formalize promises using the notion of *handled* futures, i.e., the two roles of a promise, the writing- and the reading

part, are represented by two different references, where the *handle* to the futures represents the writing-end. Apart from that, [53] are not concerned with giving an open semantics as here. On the other hand, the paper investigates the role of the heap and the reference cells, and gives a formal proof that the *only* source of non-determinism by race conditions in their language actually are the reference cells and without those, the language becomes (uniformly) confluent.⁹ Recently, an observational semantics for the (untyped) λ_{fut} -calculus has been developed in [52]. The observational equivalence is based on may- and must-program equivalence, i.e., two program fragments are considered equivalent, if, for all observing environments, they exhibit the same necessary and potential convergence behavior.

Apart from functional languages, the concept of futures has also been investigated in the object-oriented paradigm. In *Java 5*, *futures* have been introduced as part of the `java.util.concurrent` package. As *Java* does not support futures as core mechanism for parallelism, they are introduced in a library. Dereferencing of a future is done explicitly via a `get`-method (similarly to this paper). A recent paper [64] introduces *safe* futures for *Java*. The safe concept is intended to make futures and the related parallelism *transparent* and in this sense goes back to the origins of the concept: introducing parallelism via futures does not change the program's meaning. While straightforward and natural in a functional setting, safe futures in an object-oriented and thus state-based language such as *Java* require more considerations. The paper introduces a semantics which guarantees safe, i.e., transparent, futures by deriving restrictions on the scheduling of parallel executions and uses object versioning. The futures are introduced as an extension of Featherweight *Java (FJ)* [37], a core object calculus, and is implemented on top of *Jikes RVM* [10, 15]. Pratikakis et. al. [55] present a constraint-based static analysis for (transparent) futures and proxies in *Java*, based on type qualifiers and qualifier inference [32]. Also this analysis is formulated as an extension of *FJ* by type qualifiers. Similarly, Caromel et. al. [20][19][18] tackle the problem to provide *confluent*, i.e., effectively deterministic system behavior for a concurrent object calculus with futures (asynchronous sequential processes, *ASP*, an extension of Abadi and Cardelli's imperative, untyped object calculus $\text{imp}\lambda$ [1]) and in the presence of imperative features. The *ASP* model is implemented in the *ProActive Java*-library [21]. The fact, that *ASP* is derived from some (sequential, imperative) object-calculus, as in the formalization here, is more a superficial or formal similarity, in particular when being interested in the interface behavior of concurrently running objects, where the inner workings are hidden anyway. Apart from that there are some similarities and a number of differences between the work presented here and *ASP*. First of all, both calculi are centered around the notion of first-class futures, yielding active objects. The treatment, however, of getting the value back, is done differently in [18]. Whereas here, the client must explicitly claim a return value of an asynchronous method, if interested in the result, the treatment of the future references is done *implicitly* in *ASP*, i.e., the client blocks if he performs a strict operation on the future (without explicit syntax to claim the value). Apart from that, the object model is more sophisticated, in that the calculus distinguishes between active and passive objects. Here, we simple have objects, which can behave actively or passively (reactively), depending on the way they are used.

⁹Uniform confluence is a strengthening of the more well-known notion of (just ordinary) confluence; it corresponds to the diamond property of the *one-step* reduction property. For standard reduction strategies of a purely functional λ -calculus, only confluence holds, but not uniform confluence. However, the non-trivial "diamonds" in the operational semantics of λ_{fut} are caused not by different redices within one λ -term (representing one thread), but by redices from different threads running in parallel, where the reduction strategy per thread is deterministic (as in our setting, as well).

In *ASP*, the units of concurrency are the explicitly activated active objects, and each passive one is owned and belongs to exactly one active one. Especially, passive objects do not directly communicate with each other across the boundaries of concurrent activity, all such communication between concurrent activities is mediated and served by the active objects.

Related to that, a core feature of *ASP*, not present here, is the necessity to specify (also) the *receptive behavior* of the active object, i.e., in which order it is willing to process or *serve* incoming messages. The simplest serve strategy would be the willingness to accept all messages and treat them in a first-come, first-serve manner, i.e., a input-enabled FIFO strategy on the input message queue. The so-called *serve-method* is the dedicated activity of an active object to accept and schedule incoming method calls. Typically, as for instance in the FIFO case, it is given an non-terminating process, but it might also terminate, in which case the active object together with the passive objects it governs, becomes superfluous: an active object which does no service any longer does not become a passive data structure, but can no longer react in any way.

As extension of the core *ASP* calculus, [18, Chapter 10] treats *delegation* that bears some similarities with the promises here. By executing the construct $delegate(o.l(\vec{v}))$ (using our notational conventions), a thread n hands over the permission and obligation to provide eventually a value for the future reference n to method l of object o , thereby losing that permission itself. That corresponds to executing $bind\ o.l(\vec{v}) : T \hookrightarrow n$. Whereas in our setting, we must use a yet-unfulfilled promise n for that purpose, the delegation operator in *ASP* just (re-)uses the current future for that. Consequently, *ASP* does not allow the creation of promises independently from the implicit creation when asynchronously calling a method, as we do with the *promise* T construct. In this sense, the promises here are more general, as they allow to profit from delegation and have the promise as first-class entity, i.e., the programmer can pass it around as argument of methods, This, on the other hand, requires a more elaborate type system to avoid write errors on promises. This kind of error, fulfilling a promise twice, is avoided in the delegate-construct of *ASP* not by a type system, but by construction, in that the delegate-construct must be used only at the end of a method, so that the delegating activity cannot write to the future/promise after it has delegated the permission to another activity.

Further uses of futures for *Java* are reported in [49] [44] [56] [62] [61]. Futures are also integral part of *Io* [38] and *Scoop* (simple concurrent object-oriented programming) [24] [12] [50], a concurrent extension of *Eiffel*. Both languages are based on the active objects paradigm.

Benton et. al. [14] present polyphonic $C^\#$, adding concurrency to $C^\#$, featuring asynchronous methods and based on the join calculus [33] [34]. Polyphonic $C^\#$ allows methods to be declared as being asynchronous using the `async` keyword for the method type declaration. Besides that, polyphonic $C^\#$ supports so-called *chords* as synchronization or join pattern. With similar goals, *Java* has been extended by join patterns in [39] [40].

In the context of *Creol*, de Boer et. al. [27] present a formal, operational semantics for the language and extend it by *futures* (but not promises). Besides the fact, that both operational semantics ultimately formalize a comparable set of features, there are, at a technical level, a number of differences. For once, here, we simplified the language slightly mainly in two respects (apart from making it more expressive in adding promises, of course). We left out the “interleaving” operators \parallel and $\parallel\!\!\parallel$ of [27] which allows the user to express interleaving concurrency *within* one method body. Being interested in the observable interface behavior, those operations are a matter of internal,




Figure 3: Claiming a future (busy wait)

hidden behavior, namely leading to non-deterministic behavior at the interface. Since objects react non-deterministically anyhow, namely due to race conditions present independently of \parallel and /// , those operators have no impact on the possible traces at their interface. The operators might be useful as abstractions for the programmer, but without relevance for the interface traces, and so we ignored them here. Another simplification, this time influencing the interface behavior, is how the programmer can claim the value of a future. This influences, as said, the interface behavior, since the component may fetch the value of a future being part of the environment, or vice versa. Now, the design of the *Creol*-calculus in [27] is more liberal wrt. what the user is allowed to do with future references. In this paper, the interaction is rather restricted: if the client requests the value using the *claim*-operation, there are basically only two reactions. If the future computation has already been completed, the value is fetched and the client continues; otherwise it blocks until, if ever, the value is available. The bottom line is, that the client, being blocked, can never *observe* that the value is yet absent. The calculus of [27], in contrast, permits the user to *poll* the future reference directly, which gives the freedom to decide, *not* to wait for the value if not yet available. Incorporating such a construct into the language makes the *absence* of the value for a future reference observable and would complicate the behavioral interface semantics to some extent. This is also corroborated by the circumstance that the expressive power of explicit polling quite complicates the proof theory of [27] (see also the discussion in the conclusion of [27]). This is not a coincidence, since one crux of the complete Hoare-style proof systems such as in [27] is to internalize the (ideally observable) behavior into the program state by so-called auxiliary variable. In particular recording the past interaction behavior in so-called history variables is, of course, an internalization of the interface behavior, making it visible to the Hoare-assertions. As a further indication that allowing to poll a future quite adds expressivity to the language is the observation that adding a poll-operation to *ASP*, destroys a central property of *ASP*, namely confluence, as is discussed in [18, Chapter 11].

Apart from that, the combination of claiming a futures, the possibility of polling a future, and a general await-statement complicates the semantics of claiming a future: in [27], this is done by *busy-waiting*, which in practice one intends to avoid. So instead of the behavior described in Figure 1, the formalization in [27] behaves as sketched in Figure 3.

After an unsuccessful try to obtain a value of future, the requesting thread is suspended and loses the lock. In order to continue executing, the blocked thread needs two resources: the value of the future, once it is there, plus the lock again. The difference of the treatment in Figure 1 and the one of Figure 3 for [27] is the order in which the requesting thread attempts to get hold of these two resources: our formalization first

check availability of the future and afterwards re-gains the lock to continue, whereas [27] do it vice versa, leading to busy wait. The reason why it is sound to copy the future value into the local state space without already having the lock again (Figure 1) is, of course, that, once arrive, the future value remains stable and available.

In addition, our work differs also technically in the way, the operational semantics is represented. [27] formulated the (internal) operational semantics using evaluation contexts (as do, e.g., [53] for λ_{fut}), whereas we rely on a “reduction-style” semantics, making use of an appropriate notion of structural congruence. While largely a matter of taste, it seems to us that, especially in the presence of complicated synchronization mechanisms, for instance the ready queue representation of [27], the evaluation contexts do not give rise to an immediately more elegant specification of the reduction behavior. Admittedly, we ignored here the internal interleaving operators \parallel and $\parallel\parallel$, which quite contribute to the complexity of the evaluation contexts. Another technical difference, if you wish, concerns the way, the futures, threads, and objects are *represented* in the operational semantics, i.e., in the run-time syntax of the calculus. Different from our representation, the semantics makes the active-objects paradigm of *Creol* more visible, in that the activities as part of the object, more precisely, an object contains, besides the instance state, an explicit representation of the current activity (there called “process”) executing “inside” the object plus a representation of the ready-queue containing all the activities, which have been *suspended* during their execution inside the object. The scheduling between the different activities is then done by juggling them in and out of the ready-queue at the processor release points. Here, in contrast, we base our semantics on a separate representation of the involved semantics concepts: 1) classes as *generators* of objects, 2) objects carrying in the instance variables the persistent *state* of the program, thus basically forming the heap, and 3), the *parallel* activities in the form of threads. While this representation makes arguably the active-object paradigm less visible in the semantics, it on the other hand separates the concepts in a clean way, and instead of an explicit local scheduler inside the objects, the access to a share instance states of the objects is regulated by a simple, binary lock per object. So, instead of having to levels of parallelism—locally inside the objects and inter-object parallelism—the formalization achieves the same with just one conceptual level, namely: parallelism is between threads (and the necessary synchronization is done via object-locks). Additionally, our semantics is rather close to the object-calculi semantics for multi-threading as in *Java* (for instance as in [41] [42] or [59]). This allows to see the differences and similarities between the different models of concurrency, and the largely similar representation could allow are more formal comparison between the interface behaviors in the two settings.

The language *Cool* [22] [23] (concurrent, object-oriented language) is defined as an extension of C^{++} [60] for task-level parallelism on shared memory multi-processors. Concurrent execution in *Cool* is expressed by the invocation of parallel functions executing *asynchronously*. Unlike the work presented here, *Cool* future types, which correspond to the types of the form $[T]$. Further languages supporting futures include ACT-1 [46] [47], concurrent *Smalltalk* [65] [69], *Cool* [22] [23] (concurrent, object-oriented language) as a parallel extension of C^{++} [60], and of course the influential actor model [8, 35, 7], *ABCL/1*[66] [67] (in particular the extension *ABCL/f*[63]).

We have characterized the behavioral semantics of open systems, similarly to the one presented here for futures and promises, in earlier papers, especially for object-oriented languages based on *Java*-like multi-threading and synchronous method calls, as in *Java* or $C^\#$. The work [5] deals with thread classes and [4] with re-entrant monitors. In [59] the proofs of full abstraction for the sequential and multi-threaded

cases of a class-based object-calculus can be found. Poetzsch-Heffter and Schäfer [54] present a behavioral interface semantics for a class-based object-oriented calculus, however without concurrency. The language, on the other hand, features an ownership-structured heap.

Future work

An obvious way to proceed is to consider more features of the *Creol*-language, in particular inheritance and subtyping. Incorporating inheritance is challenging, as it renders the system open wrt. a new form of interaction, namely the environment inheriting behavior from a set of component classes or vice versa. Also *Creol*'s mechanisms for dynamic class upgrades should be considered from a behavioral point of view (that we expect to be quite more challenging than dealing with inheritance). An observational, black-box description of the system behavior is necessary for the compositional account of the system behavior. Indeed, the legal interface description is only a first, but necessary, step in the direction of a compositional and ultimately fully-abstract semantics, for instance along the lines of [59]. Based on the interaction trace, it will be useful to develop a logic better suited for specifying the desired interface behavior of a component than enumerating allowed traces. Another direction is to use the results in the design of a black-box testing framework, as we started for *Java* in [26]. We expect that, with the theory at hand, it should be straightforward to adapt the implementation to other frameworks featuring futures, for instance, to the future libraries of *Java 5*.

References

- [1] M. Abadi and L. Cardelli. *A Theory of Objects*. Monographs in Computer Science. Springer-Verlag, 1996.
- [2] E. Ábrahám, I. Grabe, A. Grüner, and M. Steffen. Behavioral interface description of an object-oriented language with futures and promises. Technical Report 364, University of Oslo, Dept. of Computer Science, Oct. 2007. In preparation.
- [3] E. Ábrahám, A. Grüner, and M. Steffen. Abstract interface behavior of object-oriented languages with monitors. In R. Gorrieri and H. Wehrheim, editors, *FMOODS '06*, volume 4037 of *Lecture Notes in Computer Science*, pages 218–232. Springer-Verlag, 2006.
- [4] E. Ábrahám, A. Grüner, and M. Steffen. Abstract interface behavior of object-oriented languages with monitors. *Theory of Computing Systems*, 2007. Accepted for publication. This is an extended version of the FMOODS'06 conference contribution.
- [5] E. Ábrahám, A. Grüner, and M. Steffen. Heap-abstraction for open, object-oriented systems with thread classes. *Journal of Software and Systems Modelling (SoSyM)*, 2007. This is a reworked version of the Institut für Informatik und Praktische Mathematik, Christian-Albrechts-Universität zu Kiel technical report nr. 0601 and an extended version of the CiE'06 extended abstract.
- [6] ACM. *Object Oriented Programming: Systems, Languages, and Applications (OOPSLA) '99*, 1999. In *SIGPLAN Notices*.
- [7] G. A. Agha. *Actors: a Model of Concurrent Computation in Distributed Systems*. MIT Press, Cambridge, MA, 1986.
- [8] G. A. Agha, I. A. Mason, S. F. Smith, and C. L. Talcott. Towards a theory of actor computation (extended abstract). In R. Cleaveland, editor, *Proceedings of CONCUR '92*, volume 630 of *Lecture Notes in Computer Science*, pages 565–579. Springer-Verlag, 1992.
- [9] Alice project home page. www.ps-uni-sb.de/alice, 2006.

- [10] B. Alpern, C. R. Attanasio, J. J. Barton, A. Cocchi, S. F. Hummel, D. Lieber, T. Ngo, M. Mergen, J. C. Sheperd, and S. Smith. Implementing Jalapeno in Java. In *OOPSLA'99* [6], pages 313–324. In *SIGPLAN Notices*.
- [11] P. America. Issues in the design of a parallel object-oriented language. *Formal Aspects of Computing*, 1(4):366–411, 1989.
- [12] V. Arslan, P. Eugster, P. Nienaltowski, and S. Vaucouleur. Scoop — concurrency made easy. In J. Kohlas, B. Meyer, and A. Schiper, editors, *Research Results of the DICS Program*, volume 4028 of *Lecture Notes in Computer Science*, pages 82–102. Springer, 2006.
- [13] H. Baker and C. Hewitt. The incremental garbage collection of processes. *ACM Sigplan Notices*, 12:55–59, 1977.
- [14] N. Benton, L. Cardelli, and C. Fournet. Modern concurrency abstraction for C#. *ACM Transactions on Programming Languages and Systems. Special Issue with papers from FOOL 9*, 2003.
- [15] M. G. Burke, J.-D. Choi, S. F. Fink, D. Grove, M. Hind, V. Sarkar, M. Serranon, V. C. Sreedhar, H. Srinivasan, and J. Whaley. The Jalapeno dynamic optimizing compiler. In *Proceedings of the ACM Java Grande Conference, San Francisco*, pages 129–141, 1999.
- [16] D. Caromel. Service, asynchrony and wait-by-necessity. *Journal of Object-Oriented Programming*, 2(4):12–22, Nov. 1990.
- [17] D. Caromel. Towards a method of object-oriented concurrent programming. *Communications of the ACM*, 36(9):90–102, Sept. 1993.
- [18] D. Caromel and L. Henrio. *A Theory of Distributed Objects. Asynchrony — Mobility — Groups — Components*. Springer-Verlag, 2005.
- [19] D. Caromel, L. Henrio, and B. P. Serpette. Asynchronous sequential processes. Research Report RR-4753 (version 2), INRIA Sophia-Antipolis, May 2003.
- [20] D. Caromel, L. Henrio, and B. P. Serpette. Asynchronous and deterministic objects. In *Proceedings of POPL '04*. ACM, Jan. 2004.
- [21] D. Caromel, W. Klauser, and J. Vayssière. Towards seamless computing and metacomputing in Java. *Concurrency, Practice and Experience*, 10(11-13):1043–1061, 1998. ProActive available at www.infria.fr/oasis/proactive.
- [22] R. Chandra. *The COOL Parallel Programming Language: Design, Implementation, and Performance*. PhD thesis, Stanford University, Apr. 1995.
- [23] R. Chandra, A. Gupta, and J. L. Hennessy. COOL: A language for parallel programming. In *Proceedings of the 2nd Workshop on Programming Languages and Compilers for Parallel Computing*. IEEE CS, 1989.
- [24] M. J. Compton. SCOOP: An investigation of concurrency in Eiffel. Master's thesis, Department of Computer Science, The Australian National University, 2000.
- [25] The Creol language. <http://www.ifi.uio.no/~creol>, 2007.
- [26] F. S. de Boer, M. M. Bonsangue, A. Grüner, and M. Steffen. Test driver generation from object-oriented interaction traces. In *Proceedings of the 19th Nordic Workshop of Programming Theory (NWPT'07)*, 2007.
- [27] F. S. de Boer, D. Clarke, and E. B. Johnsen. A complete guide to the future. In R. de Nicola, editor, *Proceedings of Programming Languages and Systems, 16th European Symposium on Programming, ESOP 2007, Vienna, Austria.*, volume 4421 of *Lecture Notes in Computer Science*, pages 316–330. Springer-Verlag, 2007.
- [28] The E language. www.erights.org, 2007.
- [29] C. Flanagan and M. Felleisen. The semantics of future. Technical Report TR94-238, Department of Computer Science, Rice University, 1994.

- [30] C. Flanagan and M. Felleisen. Well-founded touch optimization of parallel scheme. Technical Report TR94-239, Department of Computer Science, Rice University, 1994.
- [31] C. Flanagan and M. Felleisen. The semantics of future and an application. *Journal of Functional Programming*, 9(1):1–31, 1999.
- [32] J. Foster, M. Fändrich, and A. Aiken. A theory of type qualifiers. In *ACM Conference on Programming Language Design and Implementation*, pages 192–203. ACM, May 1999.
- [33] C. Fournet and G. Gonthier. The reflexive chemical abstract machine and the join-calculus. In *Proceedings of POPL '96*, pages 372–385. ACM, Jan. 1996.
- [34] C. Fournet and G. Gonthier. Th join calculus: A language for distributed mobile programming. In G. Barthe, P. Dybjer, L. Pinto, and J. Saraiva, editors, *APPSEM 2000*, volume 2395, 2002.
- [35] I. A. M. Gul A. Agha, S. F. Smith, and C. L. Talcott. A foundation for actor computation. *Journal of Functional Programming*, 7(1), Jan. 1997.
- [36] R. H. Halstead, Jr. Multilisp: A language for concurrent symbolic computation. *ACM Transactions on Programming Languages and Systems*, 7(4):501–538, Oct. 1985.
- [37] A. Igarashi, B. C. Pierce, and P. Wadler. Featherweight Java: A minimal core calculus for Java and GJ. In *OOPSLA'99* [6], pages 132–146. In *SIGPLAN Notices*.
- [38] Io: A small programming language. www.iolanguage.com, 2007.
- [39] G. S. Itzstein and D. Kearney. Join Java: An alternative concurrency semantics for java. Technical Report ACRC-01-001, University of South Australia, 2001.
- [40] G. S. Itzstein and D. Kearney. Applications of join Java. In *Proceedings of the Seventh Asia-Pacific Computer Systems Architectures Conference (ACSAC 2002)*, 2002.
- [41] A. Jeffrey and J. Rathke. A fully abstract may testing semantics for concurrent objects. In *Proceedings of LICS '02*. IEEE, Computer Society Press, July 2002.
- [42] A. Jeffrey and J. Rathke. Java Jr.: A fully abstract trace semantics for a core Java language. In M. Sagiv, editor, *Proceedings of ESOP 2005*, volume 3444 of *Lecture Notes in Computer Science*, pages 423–438. Springer-Verlag, 2005.
- [43] E. B. Johnsen, O. Owe, and I. C. Yu. Creol: A type-safe object-oriented model for distributed concurrent systems. *Theoretical Computer Science*, 365(1–2):23–66, Nov. 2006.
- [44] JSR 166: Concurrency utilities. www.jcp.org/en/jsr/detail?id=166, 2007.
- [45] L. Kornstaedt. Alice in the land of Oz – an interoperability-based implementation of a functional language on top of a relational language. In *Proceedings of the First Workshop on Multi-Language Infrastructure and Interoperability (BABEL'01)*, Electronic Notes in Theoretical Computer Science, Sept. 2001.
- [46] H. Liebermann. A preview of ACT-1. AI-Memo AIM-625, Artificial Intelligence Laboratory, MIT, 1981.
- [47] H. Liebermann. Concurrent object-oriented programming in ACT1. In Yonezawa and Tokoro [68].
- [48] B. Liskov and L. Shriru. Promises: Linguistic support for efficient asynchronous procedure calls in distributed systems. *SIGPLAN Notices*, 23(7):260–267, 1988.
- [49] D. A. Manolescu. Workflow enactment with continuation and future objects. In *Object Oriented Programming: Systems, Languages, and Applications (OOPSLA) '02 (Seattle, USA)*, pages 40–51. ACM, Nov. 2002. In *SIGPLAN Notices*.
- [50] B. Meyer. Systematic concurrent object-oriented programming. *Communications of the ACM*, 36(9):56–80, 1993.
- [51] L. Moreau. The semantics of scheme with future. In *International Conference on Functional Programming*, pages 146–156. ACM Press, 1996.

- [52] J. Niehren, D. Sabel, M. Schmidt-Schauß, and J. Schwinghammer. Observational semantics for a concurrent lambda calculus with reference cells and futures. *Electronic Notes in Theoretical Computer Science*, 2007.
- [53] J. Niehren, J. Schwinghammer, and G. Smolka. A concurrent lambda-calculus with futures. *Theoretical Computer Science*, 2006. Preprint submitted to TCS.
- [54] A. Poetzsch-Heffter and J. Schäfer. A representation-independent behavioral semantics for object-oriented components. In M. M. Bonsangue and E. B. Johnsen, editors, *FMOODS '07*, volume 4468 of *Lecture Notes in Computer Science*. Springer-Verlag, June 2007.
- [55] P. Pratikakis, J. Spacco, and M. W. Hicks. Transparent proxies for Java futures. In *Nineteenth Object Oriented Programming: Systems, Languages, and Applications (OOPSLA) '04*, pages 206–233. ACM, 2004. In *SIGPLAN Notices*.
- [56] R. R. Rajee, J. I. William, and M. Boyles. An asynchronous method incocation (ARMI) mechanism for Java. In *Proceedings of the ACM Workshop on Java for Science and Engineering Computation*, 1997.
- [57] A. Rossberg, D. L. Botland, G. Tack, T. Brunklau, and G. Smolka. Alice through the looking glass. In *Vol. 5 of Trends in Functional Programming*, chapter 6. Intellect Books, Bristol, 2006.
- [58] J. Schwinghammer. A concurrent λ -calculus with promises and futures. Diplomarbeit, Universität des Saarlandes, Feb. 2002.
- [59] M. Steffen. *Object-Connectivity and Observability for Class-Based, Object-Oriented Languages*. Habilitation thesis, Technische Fakultät der Christian-Albrechts-Universität zu Kiel, 2006. Submitted 4th. July, accepted 7. February 2007.
- [60] B. Stroustrup. *The C++ Programming Language*. Addison-Wesley, 1986.
- [61] T. Sysala and J. Janecek. Optimizing remote method invocation in Java. In *DEXA*, pages 29–35, June 2001.
- [62] É. Tanter, J. Noyé, D. Caromel, and P. Cointe. Partial behavioral reflection: Spatial and temporal reflection of reification. In *Eighteenth Object Oriented Programming: Systems, Languages, and Applications (OOPSLA) '03*, pages 27–46. ACM, 2003. In *SIGPLAN Notices*.
- [63] K. Taura, S. Matsuoka, and A. Yonezawa. ABCL/f: A future-based polymorphic typed concurrent object-oriented language — its design and implementation —. In *DIMACS workshop on Specification of Parallel Algorithms*, 1991.
- [64] A. Welc, S. Jagannathan, and A. Hosking. Safe futures in Java. In *Twentieth Object Oriented Programming: Systems, Languages, and Applications (OOPSLA) '05*. ACM, 2005. In *SIGPLAN Notices*.
- [65] Y. Yokote and M. Tokoro. Concurrent programming in concurrent SmallTalk. In Yonezawa and Tokoro [68], pages 129–158.
- [66] A. Yonezawa. *ABCL: An Object-Oriented Concurrent System*. MIT Press, 1990.
- [67] A. Yonezawa, J.-P. Briot, and E. Shibayama. Object-oriented concurrent programming in ABCL/1. In *Object Oriented Programming: Systems, Languages, and Applications (OOPSLA) '86 (Portland, Oregon)*, pages 258–268. ACM, 1986. In *SIGPLAN Notices* 21(11).
- [68] A. Yonezawa and M. Tokoro, editors. *Object-oriented Concurrent Programming*. MIT Press, 1987.
- [69] Y. Yonezawa, E. Shibayama, T. Takada, and Y. Honda. Modelling and programming in an object-oriented concurrent language ABCL/1. In Yonezawa and Tokoro [68], pages 55–89.

Index

- $\Delta \vdash C : \Theta$, 16
- $\hat{\Delta}, \hat{\Theta} \vdash o.l? : \vec{T} \rightarrow T$ (expected type of l), 17
- Δ (assumption name context), 6
- $\Delta \vdash C : \Theta$ (typing judgment), 6
- Γ (variable context), 7, 8
- $\Gamma; \Delta \vdash e : T :: \hat{\Gamma}, \hat{\Delta}$ (judgment), 7
- $[T]$ (type of reference to future of type T), 6
- Θ (commitment name context), 6
- Θ_1, Θ_2 , 7
- Ξ (assumption/commitment context), 16
- $\Xi \vdash_{\Delta} t : trace$, 31
- $\Xi \vdash_{\Theta} t : trace$, 31
- $v@l(\vec{v})$ (asynchronous method call), 5
- $bind\ o.l(\vec{v}) : T \leftrightarrow n$, 5
- $[\gamma]$ (core of label γ), 16
- $[a]$ (core of label a), 16
- $\mathbf{0}$ (empty component), 5
- γ_c (call label), 16
- γ_g (get label), 16
- \perp (free lock), 5
- \top (lock taken), 5
- $promise\ T$ (new promise), 5
- $new\ c$ (instantiation), 5
- $o[c, F, l]$ (object), 5
- $n\langle t \rangle$ (thread named n), 5
- $\vdash a$ (well-formed label), 17
- \rightsquigarrow (confluent internal step), 10
- $[T]$, 7
- $\xrightarrow{\tau}$ (internal step), 10
- \perp_c (undefined reference), 5
- v_{\perp} , 5
- λ_{fut} (λ -calculus with futures), 27
- ABCL/1, 30
- abstract syntax, 4
- actor, 30
- α -conversion, 12
- ASP, 27
- asynchronous method call, 5
- bn (free names), 16
- C (component), 5
- c (class name), 5
- $C^{\#}$, 28
- chord, 28
- communication labels, 16
- context
 - names update, 19
- Cool, 30
- core of a label, 16
- Creol, 4
- delegation
 - in ASP, 28
- f (field), 5
- Featherweight Java, 27
- fn (free names), 16
- future, 4, 5, 26
 - safe, 27
- future type, 30
- get (label), 16
- handle error, 13, 27
- instantiation
 - typing, 8
- Io, 28
- Java, 27
- join calculus, 28
- label
 - core, 16
 - well-formed, 17
- legal trace, 21
- lock, 5
- M (method suite), 5
- method suite, 5
- mobile code, 18
- monitor, 25
- Multilisp, 26
- n (name, reference to a future), 5
- name context
 - well-formed, 6
- O (methods and fields), 5
- object versioning, 27
- observational equivalence, 27

polyphonic $C^\#$, 28
promise, 4

reentrance, 20, 25
run-time syntax, 4

Scoop, 28
scope extrusion, 16
sequential composition, 5
step
 internal, 10
stop, 11
structural congruence, 11
subject reduction, 14, 15, 23
subsumption, 7, 10
subtyping, 7
suspend, 11

t (thread), 5
thread, 5
touch optimization, 26
type qualifier, 27
type system, 5
types, 5

Unit, 10
user syntax, 4

wait-by-necessity, 26
width subtyping, 7

List of Tables

1	Abstract syntax	4
2	Typing (components)	6
3	Typing	8
4	Typing	9
5	Internal steps	12
6	Structural congruence	13
7	Reduction modulo congruence	13
8	Labels	16
9	Typechecking labels	17
10	External steps	21
11	Legal traces (dual rules omitted)	22

List of Figures

1	Claiming a future	11
2	Scenarios	19
3	Claiming a future (busy wait)	29