A Functional View of Join

Gang Chen[1], Martin Odersky[2], Christoph Zenger[2], and Matthias Zenger[2]

[1] University of South Australia
[2] École Polytechnique Fédérale de Lausanne

Abstract. Join calculus, usually presented as a process calculus, is suitable as a foundation of both sequential and concurrent programming. We give a new operational semantics of join calculus, expressed as a reduction system with a single reduction rule similar to β reduction in lambda calculus. We also introduce a new Hindley/Milner style type system for join calculus. Compared to previous work, the type system gives more accurate types of composite and mutually recursive definitions. The type system's soundness is established by showing that our reduction rule keeps typings invariant. We present an algorithm for type inference and show its soundness and completeness.

1 Introduction

Foundational work on programming has been based traditionally on some variant of lambda calculus [Chu51], e.g. [Plo75,Fel88,Mog89,CF91,ORH93,AFM+95]. These calculi, ideally suited to sequential programming, are increasingly at odds with modern programs which are reactive in their interfaces and concurrent in their implementation. Process calculi such as π calculus [MPW92,Bou92,HY93] or fusion calculus [PV98] express the interaction aspects of computation well but lack abstraction when it comes to encoding functional and sequential computation. Laws of sequential programming usually can be recovered only with the aid of sophisticated type systems [PS93,KPT96] or name sortings [Ode95b].

Join calculus [FGL+96,FG96] can provide a unifying basis of call-by-value functional programming, sequential imperative programming and concurrent programming, and at the same time is simpler than previous calculi which try to model some combination of these areas [Bou89,Ode95a,Bou97]. As such it seems well suited to take over from lambda calculus as a generally agreed-upon foundation of programming. Join calculus has traditionally been introduced as a process calculus, and its operational semantics has been explained as a chemical abstract machine [BB90,FG96]. We show in this paper that join calculus can be regarded equivalently as a concurrent version of lambda calculus, with an operational semantics expressed as a reduction system. Our presentation of join calculus has a single reduction rule, similar to β_V-reduction in the call-by-value lambda calculus.

The motivations for this new view of join calculus are twofold. First, it strengthens the ties to lambda calculus and functional programming, which should help in carrying over results established in these areas. Second, it arguably brings with it important simplifications in the presentation of the calculus, which helps in explaining its concepts better and in streamlining proofs of technical results. For instance, the proof of our subject reduction theorem could be simplified considerably by basing it on a reduction system instead of a chemical abstract machine.

The second contribution of this paper is a revised and strengthened Hindley/Milner style type system for join calculus. Implicit polymorphism, common in functional languages, has been difficult to adapt to process calculi [Tur93]. Join calculus is an exception. Indeed, Fournet et al. [FLMR97] have proposed a Hindley/Milner style type system which provides a sound typing of joins. However, their type system still lacks polymorphism when compared to standard practice in functional programming. For instance, in a mutually recursive definition such as

```
even(xs) = isEmpty(xs)  || odd(tail(xs)),
odd(xs)  = ~isEmpty(xs) && even(tail(xs))
```

their type system would force `even` and `odd` to have monomorphic type. By contrast, in most typed functional programming languages such a composite definition would be encoded by the type checker as a fixed point operation in a let, and `even` and `odd` could both have the polymorphic type $\forall\alpha.$`List` $\alpha \rightarrow$ `Boolean`. The functional view of mutual recursion as a fixed point operation is incompatible with joins on the left-hand sides of definitions. However, we show how the polymorphism present in functional programs can be recovered by a new typing rule for local definitions.

The rest of this summary is structured as follows. Section 2 explains join calculus informally, by means of small example programs. Section 3 introduces the untyped join calculus as a reduction system. Section 4 presents our type system. Section 5 shows that the type system is semantically sound. Section 6 presents a type inference algorithm and shows that it is sound and complete. Section 7 concludes.

2 Background

Join calculus provides a simple kernel language in which call-by-value functional programming, sequential imperative programming, and concurrent programming can all be expressed. It can be characterized as a name-passing lambda calculus, extended with parallel composition and a simple construct for the synchronization of parallel computations. Here's an example of a one place buffer written in join calculus:

```
put(k, x) & empty()  =  k() & full(x),
get(k) & full(x)     =  k(x) & empty()
```

Two mutually recursive equations define four functions, `put`, `get`, `empty` and `full`. The `&` operator represents a *fork* (or parallel composition) on the right hand side of an equation, and a *join* on the left-hand side. Join calculus programs can be given a simple reduction semantics. In the example above, if a call to `get(l)` and a call to `full(y)` are present, then the two calls are rewritten to the term `l(y) & empty()`. A single call to either `get` or `full` would block until its partner is present.

Note that this is similar to the execution of Petri nets [Rei85], with function calls corresponding to places and equations corresponding to transitions. In both systems, a transition (equation) fires (rewrites) if all incoming places have tokens (all left-hand side calls are present). But join calculus programs have considerably more power than Petri nets since calls can carry arguments and definitions can be nested.

Note also that the example above uses continuation passing style. Instead of returning a result directly, functions pass their result to a continuation argument, usually named `k`. Basic join calculus enforces continuation passing style, but it is possible to define direct style extensions, which are mapped to basic join calculus via a continuation passing transform.

If we leave out the `&` operator, join calculus can be seen as a minimal functional language. For instance, here is the definition of function composition in join calculus:

```
compose(f, g, k)  =  def
                         h(x, k2) = def
                                        k1(y) = f(y, k2)
                                    in  g(x, k1)
                     in  k(h)
```

If we don't insist on continuation passing style, the function can be written in a more concise and familiar way as follows:

```
compose(f, g)  =  def h(x) = f(g(x)) in h
```

Finally, join calculus can also be seen as a basis for imperative programming, since a mutable reference can be encoded by the following two definitions of a reader function `get` and a writer function `put`, coupled via an internal function `state`.

```
get(k) & state(x)     =  k(x) & state(x),
put(y, k) & state(x)  =  k() & state(y)
```

In [FLMR97] Fournet et al. have shown that the close correspondence of join calculus with functional languages extends to type systems. They have adapted the popular Hindley/Milner system [Mil78] to join calculus, providing a sound typing of effects. Their idea is that in a definition where several functions are defined only type variables that appear in the type of a single function can be generalized. This restriction prevents type variables that appear in the types of two or more functions in a join pattern from being generalized. Such a generalization would be unsound since it would break the coupling of the types of the co-defined functions. Their approach provides a sound typing of references via their encoding into join calculus. No special treatment like weak type variables or effect types is necessary. Note however that join calculus provides a trivial form of value polymorphism [Wri95], since defined names are always functions.

For instance, in the example above the most general types for `get`, `state` and `put` are:

$$\begin{array}{ll} \texttt{get} & : (\alpha \to \diamond) \to \diamond \\ \texttt{put} & : (\alpha, () \to \diamond) \to \diamond \\ \texttt{state} & : \alpha \to \diamond \end{array}$$

Here, $\tau \to \diamond$ is the type of a function with parameter τ which does not return a result. None of these functions is polymorphic since each function shares its type variable α with all others.

This is all as it should be. However, the Fournet et al. system suffers from an annoying shortcoming: Functions whose types share type variables because one function calls the other are treated in the same way as participants of join patterns, and therefore cannot be polymorphic. For instance, in

```
f(x, k) = g(x, k),
g(x, k) = k(x)
```

neither `f` nor `g` could be given their polymorphic ML-type $\forall\alpha.(\alpha, \alpha \to \diamond) \to \diamond$. One could imagine that the problem can be side-stepped by rearranging definitions after a dependency analysis. In our example, the definitions of `g` and `f` could each in turn be checked and generalized individually, thereby eliminating the unwanted coupling of type variables. But this approach breaks down when `f` and `g` are mutually recursive, since in this case neither can be checked without the other. The usual ML encoding of mutual recursion via a fixed point combinator is not generally applicable either since such an encoding does not take join patterns into account. Fournet et al. thus provide a sound type system for join patterns at the expense of an overly monomorphic typing of composite definitions.

In the following sections, we present a type system of join calculus that allows mutually recursive functions to be polymorphic. Its main novelty is a new rule for local definitions which regards generalization of bindings as a fixed point

operation. Soundness of the new type system is established by giving an untyped operational semantics with structural equivalence rules and a β reduction rule, and showing that types are invariant under both equivalence and reduction. We also present a type inference algorithm for our system and show its soundness and completeness.

3 The Untyped Core Calculus

In this section we present our version of the untyped join calculus. The main difference compared to the original treatment [FGL$^+$96,FG96] is that we express the operational semantics as a reduction system rather than as a chemical abstract machine. This streamlines both the presentation and the proofs. The syntax of our core calculus is given as follows.

Names	$a, b, c, \ldots, x, y, z$
Terms	$M = \mathbf{def}\ D\ \mathbf{in}\ M \mid x(\overline{y}) \mid M\ \&\ M$
Definitions	$D = L = M \mid D, D \mid \epsilon$
Left-hand sides	$L = x_1(\overline{y}_1)\ \&\ \ldots\ \&\ x_n(\overline{y}_n) \qquad (n \geq 1)$

Meta-Syntactical Convention. For any expression X, we use $\overline{X}$ to indicate a sequence $X_1, \ldots, X_n$ of X's, where $n \geq 0$.

As terms M we have function calls $x(\overline{y})$, parallel composition expressed using the & operator, and local definitions, $\mathbf{def}\ D\ \mathbf{in}\ M$. A definition D is a possibly empty set of rewrite rules $L = M$. The left hand side of a rewrite rule is a *join pattern* $x_1(\overline{y}_1)\ \&\ \ldots\ \&\ x_n(\overline{y}_n)$, which defines a sequence of function names with their arguments. We require join patterns to be *linear*. That is, all function symbols in a join pattern must be pairwise different, and all function arguments must also be pairwise different. The right-hand side of a rewrite rule is an arbitrary term. The rewrite rules in a definition can be directly and mutually recursive.

The original treatment of join calculus used (|) for parallel composition and joins. We have changed this to (&) since joins and parallel composition represent conjunctions of events, whereas (|) is associated with logical or.

Definition. The set of *defined names* $\mathrm{dn}(L)$ and the set of *local names* $\mathrm{ln}(L)$ of a left-hand side L are defined as follows.

$$\mathrm{dn}(x_1(\overline{y}_1)\ \&\ \ldots\ \&\ x_n(\overline{y}_n)) = \{x_1, \ldots, x_n\}$$
$$\mathrm{ln}(x_1(\overline{y}_1)\ \&\ \ldots\ \&\ x_n(\overline{y}_n)) = \{\overline{y}_1\} \cup \ldots \cup \{\overline{y}_n\}$$

The sets of local and defined names $\mathrm{ln}(D)$, $\mathrm{dn}(D)$ of a definition D are the union of the corresponding sets $\mathrm{ln}(L)$, $\mathrm{dn}(L)$ of all its left hand sides L. The set of *free names* $\mathrm{fn}(M)$, $\mathrm{fn}(D)$ of a term and a definition are recursively defined as follows.

$$\begin{array}{lllll} \mathrm{fn}(\mathbf{def}\ D\ \mathbf{in}\ M) & = (\mathrm{fn}(D) \cup \mathrm{fn}(M)) \backslash \mathrm{dn}(D) & & \mathrm{fn}(L = M) & = \mathrm{fn}(M) \backslash \mathrm{ln}(L) \\ \mathrm{fn}(x(\overline{y})) & = \{x, \overline{y}\} & & \mathrm{fn}(D_1, D_2) & = \mathrm{fn}(D_1) \cup \mathrm{fn}(D_2) \\ \mathrm{fn}(M_1 \ \&\ M_2) & = \mathrm{fn}(M_1) \cup \mathrm{fn}(M_2) & & \mathrm{fn}(\epsilon) & = \emptyset \end{array}$$

All names occurring in a term M that are not free in M are called *bound* in M.

Hygiene condition. We assume that the set of free and bound names of every term we write are disjoint. This can always be achieved using a suitable α-renaming (see below).

Definition. A *renaming* is an idempotent map from names to names that is the identity except on a finite number of names. We let θ range over renamings. We write $\mathrm{dom}(\theta)$ for the set of names where θ is not the identity, and $\mathrm{codom}(\theta)$ for its image under θ.

Definition. *Structural equivalence* $\equiv$ is the smallest compatible equivalence relation between terms that satisfies the following laws.

1. α-renaming:
   $\mathbf{def}\ D\ \mathbf{in}\ M \equiv \mathbf{def}\ \theta D\ \mathbf{in}\ M \quad$ if $\mathrm{dom}(\theta) \subseteq \mathrm{ln}(D) \wedge \mathrm{codom}(\theta) \cap \mathrm{fn}(D) = \emptyset$
   $\mathbf{def}\ D\ \mathbf{in}\ M \equiv \mathbf{def}\ \theta D\ \mathbf{in}\ \theta M \quad$ if $\mathrm{dom}(\theta) \subseteq \mathrm{dn}(D) \wedge \mathrm{codom}(\theta) \cap \mathrm{fn}(D, M) = \emptyset$.
   In both equations we assume that θ is one-to-one.
2. $(\&)$ on terms is associative and commutative:
   $$M_1 \ \&\ M_2 \equiv M_2 \ \&\ M_1$$
   $$M_1 \ \&\ (M_2 \ \&\ M_3) \equiv (M_1 \ \&\ M_2) \ \&\ M_3$$
3. $(,)$ on definitions is associative and commutative with ϵ as an identity:
   $$D_1, D_2 \equiv D_2, D_1$$
   $$D_1, (D_2, D_3) \equiv (D_1, D_2), D_3$$
   $$D, \epsilon \equiv D$$
4. Scope extrusion:
   $$(\mathbf{def}\ D\ \mathbf{in}\ M_1) \ \&\ M_2 \equiv \mathbf{def}\ D\ \mathbf{in}\ (M_1 \ \&\ M_2) \quad \text{if } \mathrm{dn}(D) \cap \mathrm{fn}(M_2) = \emptyset$$

In the following, we will always identify terms that are structurally equivalent.

Definition. *Reduction* $\rightarrow$ is the smallest compatible relation that satisfies the β-reduction law:

$$\mathbf{def}\ D, L = M\ \mathbf{in}\ R[\theta L] \rightarrow \mathbf{def}\ D, L = M\ \mathbf{in}\ R[\theta M] \quad \text{if } \mathrm{dom}(\theta) \subseteq \mathrm{ln}(L)$$

Here R is a *reduction context* that determines where reductions can take place. The set of all reduction contexts is given by the following grammar:

$$R = [\,] \mid \mathbf{def}\ D\ \mathbf{in}\ R \mid R \ \&\ M \mid M \ \&\ R$$

The expression $R[M]$ denotes the reduction context R with its hole filled with M.

The possible shapes of reduction contexts determine which calls can be "executed", replacing some calls by the right-hand side of their corresponding definition. Generally, all calls can be executed except those that are part of the right hand side of some local definition. The definition of allowable redexes via a reduction context is common in sequential programming, however, the term *evaluation context* is usually employed. We have changed that name since there is nothing to evaluate in continuation passing style join calculus.

The operational semantics given here corresponds closely to the original treatment [FGL$^+$96]. Where we use term reduction of $\mathbf{def}\ D, L = M\ \mathbf{in}\ R[\theta L]$, they use reduction in a chemical abstract machine with a "molecule soup" that contains L and an environment that contains the definitions D, $L = M$ and all definitions in R. We don't need their heating and cooling rules that transfer definitions from the soup to the environment and back. Our structural equivalence is also similar to theirs, except that we don't have a law that fuses several definitions (i.e. their S-AND rule would correspond to an equivalence $\mathbf{def}\ D\ \mathbf{in}\ \mathbf{def}\ D'\ \mathbf{in}\ M\ \equiv\ \mathbf{def}\ D, D'\ \mathbf{in}\ M$ which is absent in our setting). This rule, which simplifies their chemical abstract machine reduction, is not necessary for our term reduction. Notably the fusion rule also breaks subject reduction; for this reason the original typed treatment of join calculus [FLMR97] abandons the law and introduces structured environments.

4 Type System

The syntax of types, type schemes, and environments is given as follows.

Type variables	α, β, γ
Types	$\tau\ = \alpha\ \mid\ (\tau_1, \ldots, \tau_n) \to \diamond$
Type schemes	$\sigma\ = \tau\ \mid\ \forall\alpha.\sigma$
Environments	$\Delta = x_1 : \tau_1, \ldots, x_n : \tau_n$
	$\Gamma = x_1 : \sigma_1, \ldots, x_n : \sigma_n$

Types are either type variables or function types. The result part of a function type is always $\diamond$ which stands for "undefined". As in the Hindley/Milner system we distinguish type schemes from types. Unlike types, type schemes can be polymorphic; that is they can contain universally quantified type variables. Environments Γ bind names to type schemes. We distinguish monomorphic environments Δ as a subset of environments which bind all names to types.

We will need several functions and operators on environments. The *domain* of an environment $\mathrm{dom}(\Gamma)$ is the set of identifiers defined in the environment. The

set of type variables occurring free in some type in Γ is written $\mathrm{tv}(\Gamma)$. The $(,)$ operator on environments is associative, with the empty environment as left and right identity. The operands Γ_1 and Γ_2 of a composite environment (Γ_1, Γ_2) are required to have disjoint domains. Proof trees where this requirement is violated are not considered valid. $\Gamma_1 \oplus \Gamma_2$ denotes the concatenation of the environments Γ_1 and Γ_2, where a binding in Γ_2 replaces any binding of the same name in Γ_1. $\Gamma_1 \cup \Gamma_2$ denotes map union of the two environments Γ_1 and Γ_2, where it is required that names bound in both environments are bound to the same type.

There are three forms for *typing judgements*, which apply to identifiers, definitions, and terms:

$\Gamma \vdash x : \sigma$	The identifier x has type scheme σ.
$\Gamma \vdash L = M : \Delta$	The rewrite rule $L = M$ produces environment Δ.
$\Gamma \vdash M : \diamond$	The term M is well-typed.

A typing judgement is *valid* if it is derivable by applications of the typing rules given below.

$\boxed{\Gamma \vdash x : \sigma}$

$$(\textsc{Taut})\ \ \Gamma \vdash x : \sigma \qquad (x : \sigma \in \Gamma) \qquad\qquad (\forall\mathrm{E})\ \frac{\Gamma \vdash x : \forall\alpha.\sigma}{\Gamma \vdash x : [\tau/\alpha].\sigma}$$

$\boxed{\Gamma \vdash L = M : \Delta}$

$$(\textsc{Eqn})\ \frac{\Gamma \oplus (\overline{y}_1 : \overline{\tau}_1, \ldots, \overline{y}_n : \overline{\tau}_n) \vdash M : \diamond}{\Gamma \vdash x_1(\overline{y}_1)\ \&\ \ldots\ \&\ x_n(\overline{y}_n) = M : (x_1 : (\overline{\tau}_1) \to \diamond, \ldots, x_n : (\overline{\tau}_n) \to \diamond)}$$

$\boxed{\Gamma \vdash M : \diamond}$

$$(\textsc{App})\ \frac{\Gamma \vdash x : (\tau_1, \ldots, \tau_n) \to \diamond \qquad \Gamma \vdash y_1 : \tau_1 \quad \ldots \quad \Gamma \vdash y_n : \tau_n}{\Gamma \vdash x(y_1, \ldots, y_n) : \diamond}$$

$$(\textsc{Par})\ \frac{\Gamma \vdash M_1 : \diamond \qquad \Gamma \vdash M_2 : \diamond}{\Gamma \vdash M_1\ \&\ M_2 : \diamond}$$

$$(\textsc{Def})\ \frac{\begin{array}{c}\Gamma \oplus (\Delta_1 \cup \ldots \cup \Delta_n) \vdash D_i : \Delta_i \\ \Gamma \oplus (\Gamma'_1 \cup \ldots \cup \Gamma'_n) \vdash \Delta_i\ \mathbf{gen}\ \Gamma'_i \quad (i = 1, \ldots, n) \\ \Gamma \oplus (\Gamma'_1 \cup \ldots \cup \Gamma'_n) \vdash M : \diamond\end{array}}{\Gamma \vdash \mathbf{def}\, D_1, \ldots, D_n\, \mathbf{in}\, M : \diamond}$$

Rule (DEF) uses an auxiliary judgement $\Gamma \vdash \Delta \ \mathbf{gen} \ \Gamma'$ which states that under the global environment Γ, the environment Δ may be generalized to Γ'. Deduction rules for this judgement are given below.

$$(\text{REFL}) \ \ \Gamma \vdash \Delta \ \mathbf{gen} \ \Delta \qquad (\forall\text{I}) \ \frac{\Gamma \vdash \Delta \ \mathbf{gen} \ (x : \sigma, \Gamma')}{\Gamma \vdash \Delta \ \mathbf{gen} \ (x : \forall\alpha.\sigma, \Gamma')} \ (\alpha \notin \text{tv}(\Gamma, \Gamma'))$$

These rules are the equivalent of the (∀I) rule of the Hindley/Milner system. They embody the restriction of join calculus that type variables which appear in the type of more than one function of a left hand side may not be generalized.

The (DEF) rule embodies the main innovation of the type system we present. It uses a double fixed point construction to deal with possible recursion between rewrite rules D_i. First, each rewrite rule D_i is checked under the monomorphic environment created by all definitions together. Then each monomorphic environment is generalized under the polymorphic environment created by all generalizations together.

It might appear more natural to check each definition immediately under the polymorphic environment created by all definitions and generalizations. Such a rule can express polymorphic recursion, which also means that it does not admit complete type inference [Hen93]. Rule (DEF) splits the single fixed point of the polymorphic recursion rule into two successive fixed points, which can both be determined by type inference.

Examples We illustrate the expressiveness of the type system with a number of examples. Consider first the definition of a local channel:

```
receive(k) & send(x)  =  k(x)
```

This will produce the environment ($\texttt{receive} : (\alpha \to \diamond) \to \diamond, \texttt{send} : \alpha \to \diamond$). The type variable α cannot be generalized since it appears in the types of both `receive` and `send`.

Next, consider the definition of a channel creator:

```
newChannel(k')  =  def receive(k) & send(x) = k(x)
                   in  k'(receive, send)
```

Then `newChannel` has type $\forall\alpha.((\alpha \to \diamond) \to \diamond, \alpha \to \diamond) \to \diamond) \to \diamond$. The type variable α can be generalized, even though it couples the types of `receive` and `send`. But both `receive` and `send` are local and thus do not form part of the generated environment

$$\texttt{newChannel} : ((\alpha \to \diamond) \to \diamond, \alpha \to \diamond) \to \diamond) \to \diamond \ .$$

In this environment, α appears only in the binding for `newChannel`, i.e. the type variable is no longer shared between several bindings. Therefore, the binding may be generalized.

Finally, consider the following definitions:

```
f(x, k) & tick()  =  g(x, k)
g(x, k) & tick()  =  f(x, k) & k(x)
```

The (admittedly a little contrived) purpose of this program fragement is to wait for an initial call of `f(x,k)`, after which the continuation `k` is called with the provided argument `x` on every second call to `tick`. The intermediate monomorphic environment Δ produced by rule (DEF) is

$$\begin{array}{ll} \texttt{f} & : (\alpha, \alpha \to \diamond) \to \diamond \\ \texttt{g} & : (\alpha, \alpha \to \diamond) \to \diamond \\ \texttt{tick} & : () \to \diamond \end{array}$$

This is also the final environment produced by the original join calculus type system [FLMR97]. By contrast, in our type system the types of `f` and `g` can both be generalized to $\forall\alpha.(\alpha, \alpha \to \diamond) \to \diamond$.

5 Type Soundness

In this section, we sketch the proofs of two theorems which state that typings are invariant under structural equivalence and reduction.

Theorem 5.1 If $M \equiv N$ and $\Gamma \vdash M : \diamond$ then $\Gamma \vdash N : \diamond$.

Theorem 5.2 (Subject Reduction) If $M \to N$ and $\Gamma \vdash M : \diamond$ then $\Gamma \vdash N : \diamond$.

Subject reduction is as far as we can get in characterizing type soundness for join calculus. It says that any reduction from a well-typed term gives a sequence of well-typed terms. For sequential languages, one usually adds the requirement that every typable term which cannot be further reduced represents an answer. In a concurrent language this requirement would imply that deadlocks are impossible for typed programs. But our type system is not strong enough to guarantee absence of deadlocks. Nevertheles, the type system is strong enough to guarantee the absence of certain kinds of run-time errors. For example, one can guarantee that function calls with a wrong number of arguments are impossible.

To prove invariance of typings under structural equivalence and reduction, we need two preliminary lemmas, which can both be shown by standard techniques.

Lemma 5.3 (Context) For all environments Γ:

$$\forall C.\, \exists \Gamma'.\, \forall M.\, \Gamma \vdash C[M] : \diamond \iff \Gamma, \Gamma' \vdash M : \diamond$$

Lemma 5.4 (Weakening) For all environments Γ, Γ' and terms M:

$$\Gamma \vdash M : \diamond \Rightarrow \Gamma, \Gamma' \vdash M : \diamond$$

provided (Γ, Γ') is well-defined, i.e. $\mathrm{dom}(\Gamma) \cap \mathrm{dom}(\Gamma') = \emptyset$.

To show Theorem 5.1, one shows that all laws of structural invariance maintain typability. This is non-trivial only for the scope extrusion law, which requires an application of the weakening lemma. One then appeals to the context lemma to show that the invariance also applies to the compatible closure of $(\equiv)$.

The key step in the proof of subject reduction (Theorem 5.2) is to show that

$$\Gamma \vdash \mathbf{def}\, D, L = M \,\mathbf{in}\, R[\theta L] : \diamond \;\Rightarrow\; \Gamma \vdash \mathbf{def}\, D, L = M \,\mathbf{in}\, R[\theta M] : \diamond \quad (1)$$

To show implication (1), we introduce two auxiliary judgements that associate definitions and contexts with environments. The first judgement, of the form $\Gamma \vdash D :_{\mathrm{gen}} \Gamma'$, associates a definition with the environment it is given in rule (Def'):

$$(\textsc{Def'})\ \frac{\begin{array}{c}\Gamma \oplus (\Delta_1 \cup \ldots \cup \Delta_n) \vdash D_i : \Delta_i \\ \Gamma \oplus (\Gamma'_1 \cup \ldots \cup \Gamma'_n) \vdash \Delta_i \ \mathbf{gen}\ \Gamma'_i \quad (i = 1, \ldots, n)\end{array}}{\Gamma \vdash (D_1, \ldots, D_n) :_{\mathrm{gen}} (\Gamma'_1 \cup \ldots \cup \Gamma'_n)}$$

The second judgement, of the form $\Gamma \vdash R :_{\mathrm{gen}} \Gamma_R$, associates a reduction context with the environment represented by all definitions which are visible at the hole of the context. This judgement is derived from the following rules.

$$\Gamma \vdash [\,] :_{\mathrm{gen}} \epsilon \qquad \frac{\Gamma \vdash D :_{\mathrm{gen}} \Gamma' \qquad \Gamma, \Gamma' \vdash R :_{\mathrm{gen}} \Gamma''}{\Gamma \vdash \mathbf{def}\, D \,\mathbf{in}\, R :_{\mathrm{gen}} \Gamma', \Gamma''}$$

$$\frac{\Gamma \vdash R :_{\mathrm{gen}} \Gamma' \qquad \Gamma \vdash M : \diamond}{\Gamma \vdash R \& M :_{\mathrm{gen}} \Gamma'}$$

Observe that environment formation with $(,)$ is always well-defined since we assume terms to be hygienic, so there is no overlap in the environments' domain.

The key property of the above judgements is:

Lemma 5.5

$$\Gamma \vdash R[M] : \diamond \Rightarrow \exists \Gamma_R.\ \Gamma \vdash R :_{\mathrm{gen}} \Gamma_R \wedge \Gamma, \Gamma_R \vdash M : \diamond$$
$$\Gamma \vdash R :_{\mathrm{gen}} \Gamma_R \wedge \Gamma, \Gamma_R \vdash M : \diamond \Rightarrow \Gamma \vdash R[M] : \diamond$$

Now, to show implication (1), assume its left-hand side is valid and pick an environment Γ_D such that $\Gamma \vdash (D, L = M) :_{\text{gen}} \Gamma_D$ and $\Gamma, \Gamma_D \vdash R[\theta L] : \diamond$. An application of Lemma 5.5 with $R \equiv \mathbf{def}\, D, L = M\, \mathbf{in}\, [\,]$ guarantees that such an environment exists. Next, pick an environment Γ_R such that $\Gamma, \Gamma_D \vdash R :_{\text{gen}} \Gamma_R$ and $\Gamma, \Gamma_D, \Gamma_R \vdash \theta L : \diamond$. Again Lemma 5.5 guarantees that such an environment exists. Next, we show that θM is typable under the same environment, that is, $\Gamma, \Gamma_D, \Gamma_R \vdash \theta M : \diamond$. Finally, we obtain $\Gamma \vdash \mathbf{def}\, D, L = M\, \mathbf{in}\, R[\theta M] : \diamond$ by reference to the second implication in Lemma 5.5. With implication (1) established, Theorem 5.2 is now shown by a simple application of the context lemma (Lemma 5.3).

6 Type Inference

In this section we present a type inference algorithm for the type system introduced before. The basic idea of type inference is adapted from the Hindley/Milner type system with the main difference lying in the rule for local definitions. Unlike in the Hindley/Milner system, definitions in Join can be composite and recursive and their left-hand sides can co-define several functions. To check a composite definition, we examine first the left-hand sides of all its equations, obtaining a monomorphic environment for every defined name. The right-hand sides are typed under this monomorphic environment. We then generalize the monomorphic environment to a possibly polymorphic one. Generalization can apply to all type variables that are not free in the enclosing environment, except for those that are present in the types of several co-defined functions. Soundness and completeness results for our type inference algorithm are given at the end of this section.

Before presenting the formalization of the type inference algorithm, we introduce some auxiliary functions in Figure 1. $inst(\sigma)$ replaces the bound type variables in type scheme σ with fresh type variables. Function $denv$ takes a left-hand side of an equation and returns a new type environment for all the defined names. With function $lenv$ it is possible to map such a type environment for a left-hand side $x_1(\overline{y}_1) \,\&\, \ldots \,\&\, x_n(\overline{y}_n)$ to the corresponding local type environment with domain $\{\overline{y}_1, \ldots, \overline{y}_n\}$. Finally, the last two functions mtv and gen are used for the generalization of monomorphic type environments. For a set of names, mtv returns all type variables that appear free in more than one of the names types. $gen(\Gamma, mono, \Delta)$ generalizes a complete type environment Δ by universally quantifying all free variables in types, that are not free in the environment Γ and that are not element of the set $mono$.

The inference algorithm is shown in Figure 2. It is an adaptation of algorithm W introduced by Damas and Milner in [DM82]. Like the former it reduces a type inference problem to a unification problem, which can be solved by a standard algorithm mgu. For different syntactic domains of our core calculus we define separate functions W_x, W_M and W_D.

$$
\begin{array}{ll}
\mathit{inst}(\forall\alpha_1 \ldots \alpha_n.\tau) & = \texttt{let fresh } \beta_1, \ldots, \beta_n \\
 & \quad \texttt{in } [\beta_1/\alpha_1, \ldots, \beta_n/\alpha_n]\tau \\
\mathit{denv}(x_1(\overline{y}_1) \,\&\, \ldots \,\&\, x_n(\overline{y}_n)) & = \texttt{let fresh } \overline{\alpha}_1, \ldots, \overline{\alpha}_n \quad (|\overline{y}_i| = |\overline{\alpha}_i|) \\
 & \quad \texttt{in } [x_1 : \overline{\alpha}_1 \to \diamond, \ldots, x_n : \overline{\alpha}_n \to \diamond] \\
\mathit{lenv}([x_1 : \overline{\tau}_1 \to \diamond, \ldots, x_n : \overline{\tau}_n \to \diamond], & \\
\quad x_1(\overline{y}_1) \,\&\, \ldots \,\&\, x_n(\overline{y}_n)) & = [\overline{y}_1 : \overline{\tau}_1, \ldots, \overline{y}_n : \overline{\tau}_n] \\
\mathit{mtv}(\Delta, V) & = \bigcup_{x,y \in V, x \neq y} \mathrm{tv}(\Delta(x)) \cap \mathrm{tv}(\Delta(y)) \\
\mathit{gen}(\Gamma, \mathit{mono}, [x_1 : \tau_1, \ldots, x_n : \tau_n]) & = [x_1 : \mathit{gen}(\Gamma, \mathit{mono}, \tau_1), \ldots, x_n : \mathit{gen}(\Gamma, \mathit{mono}, \tau_n)] \\
\mathit{gen}(\Gamma, \mathit{mono}, \tau) & = \texttt{let } \{\alpha_1, \ldots, \alpha_n\} = \mathrm{tv}(\tau) \backslash (\mathrm{tv}(\Gamma) \cup \mathit{mono}) \\
 & \quad \texttt{in } \forall\alpha_1 \ldots \alpha_n.\tau
\end{array}
$$

Fig. 1. Auxiliary functions for type inference

$$
\begin{array}{ll}
W_x & : \mathit{Sub} \times \mathit{Env} \times \mathit{Name} \to \mathit{Type} \\
W_x(\phi, \Gamma, x) & = \phi\, \mathit{inst}(\Gamma(x)) \\
 & \\
W_M & : \mathit{Sub} \times \mathit{Env} \times \mathit{Term} \to \mathit{Sub} \\
W_M(\phi, \Gamma, x\,(y_1, \ldots, y_n)\,) & = \texttt{let all } i = 1 \ldots n\text{: } \tau_i = W_x(\phi, \Gamma, y_i) \\
 & \quad \texttt{in } \mathit{mgu}(W_x(\phi, \Gamma, x) = \overline{\tau} \to \diamond)\phi \\
W_M(\phi, \Gamma, M_1 \,\&\, M_2) & = W_M(W_M(\phi, \Gamma, M_1), \Gamma, M_2) \\
W_M(\phi, \Gamma, \textbf{def } D \textbf{ in } M) & = \texttt{let } (\phi', \Gamma') = W_D(\phi, \Gamma, D) \\
 & \quad \texttt{in } W_M(\phi', \Gamma \oplus \Gamma', M) \\
 & \\
W_D & : \mathit{Sub} \times \mathit{Env} \times \mathit{Defn} \to \mathit{Sub} \times \mathit{Env} \\
W_D(\phi, \Gamma, L_1{=}M_1, \ldots, L_n{=}M_n) & = \texttt{let all } i = 1 \ldots n\text{: } \Delta_i = \mathit{denv}(L_i) \\
 & \qquad \phi_0 = \mathit{mgu}(\textstyle\bigwedge_{x \in \mathrm{dom}(\Delta_i) \cap \mathrm{dom}(\Delta_j)} \Delta_i(x) = \Delta_j(x))\phi \\
 & \qquad \Delta = \phi_0\Delta_1 \cup \ldots \cup \phi_0\Delta_n \\
 & \qquad \texttt{all } i = 1 \ldots n\text{:} \\
 & \qquad\quad \phi_i = W_M(\phi_{i-1}, \Gamma \oplus \Delta \oplus \mathit{lenv}(\Delta_i, L_i), M_i) \\
 & \qquad \mathit{mono} = \mathit{mtv}(\phi_n\Delta, \mathrm{dn}(L_1)) \cup \ldots \cup \mathit{mtv}(\phi_n\Delta, \mathrm{dn}(L_n)) \\
 & \quad \texttt{in } (\phi_n, \mathit{gen}(\phi_n\Gamma, \mathit{mono}, \phi_n\Delta))
\end{array}
$$

Fig. 2. Type inference algorithm

Type inference for names and terms is straightforward, so we focus our interest on W_D. Given a substitution ϕ and a type environment Γ, W_D returns a polymorphic type environment Γ' and a new substitution ϕ' for a tuple of definitions $D_1, \ldots, D_n$. For every equation D_i, W_D first determines environments Δ_i consisting of all defined names. Then these environments are combined into a monomorphic type environment Δ by unifying types of names that are defined in more than one Δ_i. Now, for all equations $L_i = M_i$ type inference is applied to the right-hand sides M_i under a combination of the enclosing type environment Γ, environment Δ and the type environment local to the definition D_i. Finally, the types in Δ have to be generalized. As we mentioned above, type variables cannot be generalized if they appear free in the types of several co-defined names. Therefore we have to compute the set of non generalizable variables *mono* by applying function *mtv* to all equations. Every type τ in Δ can be generalized by universally quantifying all free type variables in τ that are neither free in the enclosing environment nor element of set *mono*.

It is obvious that the type inference algorithm terminates, since the expression argument is getting smaller in each recursive call. Furthermore W is sound and complete in the following sense:

Corollary 6.1 (Soundness) $W_M(id, \Gamma, M) = \phi \;\Rightarrow\; \phi\Gamma \vdash M : \diamond$.

Corollary 6.2 (Completeness) $\psi\Gamma \vdash M : \diamond \;\Rightarrow\; W_M(id, \Gamma, M) = \phi$ and $\phi \leq_{\mathrm{tv}(\Gamma)} \psi$.

All given and resulting substitutions are assumed/claimed to be idempotent and we write $\phi \leq_U^{\psi'} \psi$ iff $\psi'\phi = \psi'$, and $\psi'|_{U \cup \mathrm{dom}(\psi)} = \psi$ where $\mathrm{dom}(\phi) = \{u | \phi u \neq u\}$. This expresses the fact that ϕ is a more general substitution than ψ on U. We write $\phi \leq_U \psi$ if there exists a ψ' such that $\phi \leq_U^{\psi'} \psi$.

Soundness includes that if M is not typable in an environment Γ, $W_M(id, \Gamma, M)$ will fail in some *mgu*-computation. A notion of principal types is hardly useful here, because the type of a term is always $\diamond$. The two above claims are special cases of these soundness and completeness theorems:

Theorem 6.3 (Soundness)

$$
\begin{aligned}
&W_x(\phi, \Gamma, x) = \tau && \Rightarrow \phi\Gamma \vdash x : \tau \text{ and } \phi\tau = \tau \\
&W_M(\phi, \Gamma, M) = \phi' && \Rightarrow \phi'\Gamma \vdash M : \diamond \\
&W_D(\phi, \Gamma, D) = (\phi', \Gamma') && \Rightarrow \phi'\Gamma \vdash D :_{\mathrm{gen}} \Gamma' \text{ and } \phi'\Gamma' = \Gamma'
\end{aligned}
$$

and $\phi' = \phi'\phi$.

Theorem 6.4 (Completeness) If $\phi \leq_U^{\chi} \psi$ then

$$\psi\Gamma \vdash x : \psi\tau \quad \Rightarrow W_x(\phi, \Gamma, x) = (V, \tau') \text{ and } \chi'\tau' = \chi\tau$$
$$\psi\Gamma \vdash M : \diamond \quad \Rightarrow W_M(\phi, \Gamma, M) = (V, \phi')$$
$$\psi\Gamma \vdash D :_{\text{gen}} \psi\Gamma' \Rightarrow W_D(\phi, \Gamma, D) = (V, \phi', \Gamma'') \text{ and } \chi'\Gamma'' \subseteq \chi\Gamma'$$

for some fresh V and $\phi' \leq_{C(V)}^{\chi'} \chi$, where $\phi' = \phi$ for the first case.

Here, a type scheme $\forall\overline{\beta}.\tau'$ is called a *generic instance* of $\forall\overline{\alpha}.\tau$ iff there exist $\overline{\tau}''$ with $[\overline{\tau}''/\overline{\alpha}]\tau = \tau'$ and $\text{tv}(\tau) \cap \overline{\beta} = \emptyset$. We write $\forall\overline{\alpha}.\tau \subseteq \forall\overline{\beta}.\tau'$. An environment Γ' is a generic instance of Γ iff $\text{dom}(\Gamma') \subseteq \text{dom}(\Gamma)$ and for each $x \in \text{dom}(\Gamma')$ we have $\Gamma(x) \subseteq \Gamma'(x)$. The relation $\subseteq$ is reflexive, transitive and stable under substitution. $C(V)$ denotes the complement of V with respect to the set of all type variables. We use the judgement $\Gamma \vdash D :_{\text{gen}} \Gamma'$, which was introduced in the section on type soundness.

The proofs of these theorems proceed on structural induction on the size of the term to type-check. The cases W_x and W_M are easily adapted from standard Hindley/Milner type systems and so is W_D but for the following generalization lemma, which reflects the new generalization rule.

Lemma 6.5 (Generalization)

$$\forall i.\ \Gamma_i' = gen(\Gamma, \bigcup_j mtv(\Delta_j, \text{dn}(L_j)), \Delta_i) \Rightarrow \forall i.\ \Gamma \oplus (\Gamma_1' \cup \ldots \cup \Gamma_n') \vdash \Delta_i \textbf{ gen } \Gamma_i'$$
$$\forall i.\ \Gamma \oplus (\Gamma_1' \cup \ldots \cup \Gamma_n') \vdash \Delta_i \textbf{ gen } \Gamma_i' \Rightarrow \forall i.\ gen(\Gamma, \bigcup_j mtv(\Delta_j, \text{dn}(L_j)), \Delta_i) \subseteq \Gamma_i'$$

For the first statement we check, that all variables, which are not in $\text{tv}(\Gamma) \cup \bigcup_j mtv(\Delta_j, \text{dn}(L_j))$ can really be generalized. The second statement is based on the observations that for arbitrary $\Gamma \vdash \Delta$ **gen** Γ' we have $\text{tv}(\Gamma) \cap \text{tv}(\Delta) \subseteq \text{tv}(\Gamma')$ and $\text{tv}(\Delta(x_i)) \cap \text{tv}(\Delta(x_j)) \subseteq \text{tv}(\Gamma'(x_i)) \cap \text{tv}(\Gamma'(x_j))$ for $i \neq j$.

The complete and detailed proofs can be found in the technical report [OZCZ99].

7 Conclusion

We have re-interpreted join calculus as a reduction system, and have studied a Hindley/Milner style type system for it which is more accurate than previous work. This work was motivated by the desire to develop a simple foundation for languages that support functional, imperative, and concurrent programming. Current examples of such languages include JoCaml [Fes98], based on original join calculus, Pict [PT97] or Piccola [ALSN99], based on π calculus, as well as CML [Rep91], Facile [GMP89], or Erlang [AVW93], which are based on different, generally more complex foundations. We are currently defining and implementing a kernel programming language which is based on our interpretation of join calculus.

References

[AFM+95] Zena M. Ariola, Matthias Felleisen, John Maraist, Martin Odersky, and Philip Wadler. A call-by-need lambda calculus. In *Proc. 22nd ACM Symposium on Principles of Programming Languages*, pages 233–246, January 1995.

[ALSN99] Franz Achermann, Markus Lumpe, Jean-Guy Schneider, and Oscar Nierstrasz. Piccola - a small composition language. Submitted for Publication, available from `http://ww.iam.unibe.ch/~scg/Research/Piccola`, 1999.

[AVW93] J. Armstrong, R. Virding, and M. Williams. *Concurrent Programming in Erlang*. Prentice Hall, 1993.

[BB90] Gérard Berry and Gérard Boudol. The chemical abstract machine. In *Proc. 17th ACM Symposium on Principles of Programming Languages*, pages 81–94, January 1990.

[Bou89] Gérard Boudol. Towards a lambda-calculus for concurrent and communicating systems. In J. Díaz and F. Orejas, editors, *Proceedings TAPSOFT '1989*, pages 149–161, New York, March 1989. Springer-Verlag. Lecture Notes in Computer Science 351.

[Bou92] Gérard Boudol. Asynchrony and the pi-calculus. Research Report 1702, INRIA, May 1992.

[Bou97] Gérard Boudol. The pi-calculus in direct style. In *Proc. 24th ACM Symposium on Principles of Programming Languages*, pages 228–241, 1997.

[CF91] Erik Crank and Matthias Felleisen. Parameter-passing and the lambda-calculus. In *Proc. 18th ACM Symposium on Principles of Programming Languages, Orlando, Florida*, pages 233–244, January 1991.

[Chu51] Alonzo Church. *The Calculi of Lambda-Conversion*, volume 6 of *Annals of Mathematics Studies*. Princeton University Press, second edition, 1951.

[DM82] Luis Damas and Robin Milner. Principal type schemes for functional programs. In *Proc. 9th ACM Symposium on Principles of Programming Languages*, pages 207–211, January 1982.

[Fel88] Matthias Felleisen. λ-v-CS: An extended λ-calculus for scheme. In *Proceedings of the 1988 ACM Conference on Lisp and Functional Programming*, pages 72–84, Snowbird, Utah, July 1988.

[Fes98] Fabrice Le Fessant. *The JoCaml reference manual*. INRIA Rocquencourt, 1998. Available from `http://join.inria.fr`.

[FG96] Cédric Fournet and Georges Gonthier. The reflexive chemical abstract machine and the join-calculus. In *Proceedings of the 23rd ACM Symposium on Principles of Programming Languages*, pages 372–385, St. Petersburg Beach, Florida, January 21-24 1996. ACM.

[FGL+96] Cédric Fournet, Georges Gonthier, Jean-Jacques Lévy, Luc Maranget, and Didier Rémy. A calculus of mobile agents. In *7th International Conference on Concurrency Theory (CONCUR'96)*, pages 406–421, Pisa, Italy, August 26-29 1996. Springer-Verlag. LNCS 1119.

[FLMR97] Cédric Fournet, Cosimo Laneve, Luc Maranget, and Didier Rémy. Implicit typing à la ML for the join-calculus. In *Proc. of the 1997 8th International Conference on Concurrency Theory*. Springer-Verlag, 1997.

[GMP89] Alessandro Giacalone, Prateek Mishra, and Sanjiva Prasad. Facile: A symmetric integration of concurrent and functional programming. *International Journal of Parallel Programming*, 18(2):121–160, April 1989.

[Hen93] Fritz Henglein. Type inference with polymorphic recursion. *ACM Transactions on Programming Languages and Systems*, 15(1):253–289, April 1993.

[HY93] Keiho Honda and Nobuko Yoshida. On reduction-based process semantics. In *Proc. 13th Conf. on Foundations of Softawre Technology and Theoretical Computer Science*, pages 373–387, December 1993.

[KPT96] Naoki Kobayashi, Benjamin C. Pierce, and David N. Turner. Linearity and the pi-calculus. In *Proc. 23rd ACM Symposium on Principles of Programming Languages*, 1996.

[Mil78] Robin Milner. A theory of type polymorphism in programming. *Journal of Computer and System Sciences*, 17:348–375, Dec 1978.

[Mog89] Eugenio Moggi. Computational lambda-calculus and monads. In *Proceedings 1989 IEEE Symposium on Logic in Computer Science*, pages 14–23. IEEE, June 1989.

[MPW92] Robin Milner, Joachim Parrow, and David Walker. A calculus of mobile processes. *Information and Computation*, 100:1–77, 1992.

[Ode95a] Martin Odersky. Applying π: Towards a basis for concurrent imperative programming. In *Proc. 2nd ACM SIGPLAN Workshop on State in Programming Languages*, January 1995.

[Ode95b] Martin Odersky. Polarized name passing. In *Proc. FST & TCS*, LNCS 1026, pages 324–337. Springer Verlag, December 1995.

[ORH93] Martin Odersky, Dan Rabin, and Paul Hudak. Call-by-name, assignment, and the lambda calculus. In *Proc. 20th ACM Symposium on Principles of Programming Languages*, pages 43–56, January 1993.

[OZCZ99] Martin Odersky, Matthias Zenger, Gang Chen, and Christoph Zenger. A functional view of join. Technical Report ACRC-99-016, University of South Australia, 1999.

[Plo75] Gordon D. Plotkin. Call-by-name, call-by-value, and the λ-calculus. *Theoretical Computer Science*, 1:125–159, 1975.

[PS93] Benjamin Pierce and Davide Sangiorgi. Typing and subtyping for mobile processes. In *IEEE Symposium on Logic in Computer Science*, June 1993.

[PT97] Benjamin C. Pierce and David N. Turner. Pict: A programming language based on the pi-calculus. Technical Report CSCI 476, Computer Science Department, Indiana University, 1997. To appear in *Proof, Language and Interaction: Essays in Honour of Robin Milner*, Gordon Plotkin, Colin Stirling, and Mads Tofte, editors, MIT Press, 1999.

[PV98] Joachim Parrow and Björn Victor. The fusion calculus. In *Proc. 13th IEEE Symposium on Logic in Computer Science*, pages 176–185. IEEE Computer Society Press, 1998.

[Rei85] W. Reisig. *Petri Nets*, volume 4 of *EATCS Monographs on Theoretical Computer Science*. Springer-Verlag, 1985.

[Rep91] John H. Reppy. CML: A higher-order concurrent language. In *Proceedings of the ACM SIGPLAN '91 Conference on Programming Language Design and Implementation*, pages 293–305, June 1991.

[Tur93] David N. Turner. *The Pi calculus: Types, Polymorphism and Implementation*. PhD thesis, LFCS, University of Edinburgh, 1993.

[Wri95] Andrew K. Wright. Simple imperative polymorphism. *Lisp and Symbolic Computation*, 8(4):343–356., December 1995.