The Geometry of Types

Ugo Dal Lago   Barbara Petit

Abstract

We show that time complexity analysis of higher-order functional programs can be effectively reduced to an arguably simpler (although computationally equivalent) verification problem, namely checking first-order inequalities for validity. This is done by giving an efficient inference algorithm for linear dependent types which, given a PCF term, produces in output both a linear dependent type and a cost expression for the term, together with a set of proof obligations. Actually, the output type judgement is derivable *iff* all proof obligations are valid. This, coupled with the already known *relative completeness* of linear dependent types, ensures that no information is lost, i.e., that there are no false positives or negatives. Moreover, the procedure reflects the difficulty of the original problem: simple PCF terms give rise to sets of proof obligations which are easy to solve. The latter can then be put in a format suitable for automatic or semi-automatic verification by external solvers. Ongoing experimental evaluation has produced encouraging results, which are briefly presented in the paper.

1   Introduction

One of the most crucial non-functional properties of programs is the amount of resources (like time, memory and power) they need when executed. Deriving upper bounds on the resource consumption of programs is crucial in many cases, but is in fact an undecidable problem as soon as the underlying programming language is non-trivial. If the units of measurement in which resources are measured become concrete and close to the physical ones, the problem becomes even more complicated, given the many transformation and optimisation layers programs are applied to before being executed. A typical example is the one of WCET techniques adopted in real-time systems [30], which do not only need to deal with how many machine instructions a program corresponds to, but also with how much time each instruction costs when executed by possibly complex architectures (including caches, pipelining, etc.), a task which is becoming even harder with the current trend towards multicores architectures.

A different approach consists in analysing the *abstract* complexity of programs. As an example, one can take the number of instructions executed by the program as a measure of its execution time. This is of course a less informative metric, which however becomes more accurate if the actual time complexity *of each instruction* is kept low. One advantage of this analysis is the independence from the specific hardware platform executing the program at hand: the latter only needs to be analysed once. A variety of verification techniques have been employed in this context, from abstract interpretation [19] to type systems [23] to program logics [22] to interactive theorem proving[1].

Among the many type-based techniques for complexity analysis, a recent proposal consists in going towards systems of *linear dependent types*, as suggested by Marco Gaboardi and the first author [12]. In linear dependent type theories, a judgement has the form $\vdash_I t : \sigma$, where σ is the type of t and I is its *cost*, an estimation of its time complexity. In this paper, we show that the problem of checking, given a PCF term t and I, whether $\vdash_I t : \sigma$ holds can be efficiently reduced to the one of checking the truth of a set of proof obligations, themselves formulated in the language of a *first-order* equational program. Interestingly, simple λ-terms give rise to simple equational programs. In other words, linear dependent types are not only a sound and *relatively* complete methodology for inferring *time* bounds of programs [12, 13]: they also allow to reduce complexity analysis to an arguably simpler (although computationally equivalent) problem which

---

[1]A detailed discussion with related work is in Section 6

Figure 1: General scheme of the type inference algorithm.

is much better studied and for which a variety of techniques and concrete tools exist [6]. Noticeably, the bounds one obtains this way translate to bounds on the number of steps performed by evaluation machines for the λ-calculus, which means that the induced metrics are not too abstract after all. The type inference algorithm is described in Section 4.

The scenario, then, becomes similar to the one in Floyd-Hoare program logics for imperative programs, where completeness holds [9] (at least for the simplest idioms [7]) *and* weakest preconditions can be generated automatically (see, e.g., [22]). A benefit of working with functional programs is that type inference — the analogue of generating WPs — can be done compositionally without the need of guessing invariants.

Linear dependent types are simple types annotated with some *index terms*, i.e. first-order terms reflecting the value of data flowing inside the program. Type inference produces in output a type derivation, a set of inequalities (which should be thought of as proof obligations) and an equational program $\mathcal{E}$ giving meaning to function symbols appearing in index terms (see Figure 1). A natural thing to do once $\mathcal{E}$ and the various proof obligations are available is to try to solve them automatically, as an example through SMT solvers. If automatically checking the inequalities for truth does not succeed (which must happen, in some cases), one can anyway find useful information in the type derivation, as it tells you precisely *which* data every symbol corresponds to. We elaborate on this issue in Section 5.

But where does linear dependency come from? Linear dependent types can be seen as a way to turn Girard's geometry of interaction [17] (or, equivalently, AJM games [1]) into a type system for the λ-calculus: the equational program one obtains as a result of type inference of a term t is nothing but as a description of a token machine for t. In presence of linear dependency, any term which can possibly be duplicated, can receive different, although uniform, types, similarly to what happens in BLL [18]. As such, this form of dependency is significantly simpler than the one of, e.g., the calculus of inductive constructions.

2 Linear Dependency at a Glance

Traditionally, type systems carry very little information about the *value* of data manipulated by programs, instead focusing on their *nature*. As an example, all (partial recursive) functions from natural numbers to natural numbers can be typed as Nat $\Rightarrow$ Nat in the λ-calculus with natural numbers and higher-order recursion, also known as PCF [27]. This is not an intrinsic limit of the type-based analysis of programs, however: much richer type disciplines have flourished in the last twenty years [21, 14, 10, 3]. All of them guarantee stronger properties for typable programs, the price being a more complicated type language and computationally more difficult type inference and checking problems. As an example, sized types [21] are a way to ensure termination of functional programs based on size information. In systems of sized types, a program like

$$t = \lambda x.\lambda y.\text{add (add } x\ y)\ (\text{succ } y)$$

can be typed as $\mathtt{Nat}_a \Rightarrow \mathtt{Nat}_b \Rightarrow \mathtt{Nat}_{a+2b+1}$, and in general as $\mathtt{Nat}_a \Rightarrow \mathtt{Nat}_b \Rightarrow \mathtt{Nat}_I$, where $I \geqslant a + 2b + 1$. In other words, the PCF type $\mathtt{Nat}$ is refined into $\mathtt{Nat}_I$ (where I is an arithmetical expression) whose semantics is the set of all natural numbers smaller or equal to I, i.e. the interval $[0, I] \subseteq \mathbb{N}$. The role of size information is to ensure that all functions terminate, and this is done by restricting the kind of functions of which one is allowed to form fixpoints. Sized types are nonlinear: arguments to functions can be freely duplicated. Moreover, the size information is only approximate, since the expression labelling base types is only an *upper* bound on the size of typable values.

Linear dependent types can be seen as a way to inject precision and linearity into sized types. Indeed, t receives the following type in $\mathsf{d\ell PCF_V}$ [13]:

$$\mathtt{Nat}[0, a] \xrightarrow{c<I} \mathtt{Nat}[0, b] \xrightarrow{d<J} \mathtt{Nat}[0, a + 2b + 1]. \tag{1}$$

As one can easily realise, $\mathtt{Nat}[K, H]$ is the type of all natural numbers in the interval $[K, H] \subseteq \mathbb{N}$. Moreover, $\sigma \xrightarrow{b<J} \tau$ is the type of *linear* functions from σ to τ which can be copied by the environment J times. The J copies of the function have types obtained by substituting $0, \dots, J-1$ for b in σ and τ. This is the key idea behind linear dependency. The type (1) is imprecise, but can be easily turned into

$$\mathtt{Nat}[a, a] \xrightarrow{c<I} \mathtt{Nat}[b, b] \xrightarrow{d<J} \mathtt{Nat}[a + 2b + 1, a + 2b + 1], \tag{2}$$

itself a type of t. In the following, the singleton interval type $\mathtt{Nat}[K, K]$ is denoted simply as $\mathtt{Nat}[K]$.

Notice that linear dependency is not exploited in (2), e.g., d does not appear free in $\mathtt{Nat}[b]$ nor in $\mathtt{Nat}[a + 2b + 1]$. Yet, (2) precisely captures the functional behaviour of t. If d does not appear free in σ nor in τ, then $\sigma \xrightarrow{d<I} \tau$ can be abbreviated as $\sigma \xrightarrow{I} \tau$. Linear dependency becomes necessary in presence of higher order functions. Consider, as another example, the term

$$u = \lambda x.\lambda y.\ \mathtt{ifz}\ y\ \mathtt{then}\ \underline{0}\ \mathtt{else}\ xy$$

u has simple type $(\mathtt{Nat} \Rightarrow \mathtt{Nat}) \Rightarrow \mathtt{Nat} \Rightarrow \mathtt{Nat}$. One way to turn it into a linear dependent type is the following

$$(\mathtt{Nat}[a] \xrightarrow{H} \mathtt{Nat}[I]) \xrightarrow{K} \mathtt{Nat}[a] \xrightarrow{L} \mathtt{Nat}[J], \tag{3}$$

where J equals 0 when $a = 0$ and J equals I otherwise. Actually, u has type (3) *for every* I and J, provided the two expressions are in the appropriate relation. Now, consider the term

$$v = (\lambda x.\lambda y.(x\ \mathrm{pred}\ (x\ \mathrm{id}\ y)))\ u.$$

The same variable x is applied to the identity id and to the predecessor pred. Which type should we give to the variable x and to u, then? If we want to preserve precision, the type should reflect *both* uses of x. The right type for u is actually the following:

$$(\mathtt{Nat}[a] \xrightarrow{1} \mathtt{Nat}[I]) \xrightarrow{c<2} \mathtt{Nat}[a] \xrightarrow{1} \mathtt{Nat}[J] \tag{4}$$

where both I and J evaluate to a if $c = 0$ and to $a - 1$ otherwise. If id is replaced by succ in the definition of v, then (4) becomes even more complicated: the first "copy" of J is fed not with a but with either 0 or $a + 1$.

Linear dependency precisely consists in allowing different copies of a term to receive types which are indexed differently (although having the same "functional skeleton") and to represent all of them in compact form. This is in contrast to, e.g., intersection types, where the many different ways a function uses its argument could even be structurally different. This, as we will see in Section 3, has important consequences on the kind of completeness results one can hope for: if the language in which *index terms* are written is sufficiently rich, then the obtained system is complete in an intensional sense: a precise type can be given to *every* terminating t *having type* $\mathtt{Nat} \Rightarrow \mathtt{Nat}$.

$$t = \lambda f.\ \mathtt{ifz}\ (f\ \underline{0})\ \mathtt{then}\ \underline{0}\ \mathtt{else}\ f(f\ \underline{0}); \qquad u = (\lambda x.\lambda y.\mathrm{add}\,x\,y)\ \underline{3}; \qquad \xi = f \mapsto \langle\,\lambda y.\mathrm{add}\,x\,y\,;\,x \mapsto \langle\,\underline{3}\,\rangle\,\rangle.$$

$$
\begin{array}{rl}
\langle\,tu\,\rangle\ \star\ \diamond\quad \succ & \langle\,t\,\rangle\ \star\ \mathsf{arg}\langle\,u\,\rangle\cdot\diamond \\
\succ & \langle\,u\,\rangle\ \star\ \mathsf{fun}\langle\,t\,\rangle\cdot\diamond \\
\succ & \langle\,\lambda x.\lambda y.\mathrm{add}\,x\,y\,\rangle\ \star\ \mathsf{arg}\langle\,\underline{3}\,\rangle\cdot\mathsf{fun}\langle\,t\,\rangle\cdot\diamond \\
\succ^* & \langle\,\lambda y.\mathrm{add}\,x\,y\,;\,x \mapsto \langle\,\underline{3}\,\rangle\,\rangle\ \star\ \mathsf{fun}\langle\,t\,\rangle\cdot\diamond \\
\succ & \langle\,\mathtt{ifz}\ (f\,\underline{0})\ \mathtt{then}\ \underline{0}\ \mathtt{else}\ f(f\,\underline{0})\,;\,\xi\,\rangle\ \star\ \diamond \\
\succ & \langle\,f\,\underline{0}\,;\,\xi\,\rangle\ \star\ \mathsf{fork}\langle\,\underline{0}\,;\,f(f\,\underline{0})\,;\,\xi\,\rangle\cdot\diamond \\
\succ & \langle\,f\,;\,\xi\,\rangle\ \star\ \mathsf{arg}\langle\,\underline{0}\,;\,\xi\,\rangle\cdot\mathsf{fork}\langle\,\underline{0}\,;\,f(f\,\underline{0})\,;\,\xi\,\rangle\cdot\diamond \\
\succ & \langle\,\lambda y.\mathrm{add}\,x\,y\,;\,x \mapsto \langle\,\underline{3}\,\rangle\,\rangle\ \star\ \mathsf{arg}\langle\,\underline{0}\,;\,\xi\,\rangle\cdot\mathsf{fork}\langle\,\underline{0}\,;\,f(f\,\underline{0})\,;\,\xi\,\rangle\cdot\diamond \\
\succ^* & \langle\,\mathrm{add}\,x\,y\,;\,(y \mapsto \langle\,\underline{0}\,;\,\xi\,\rangle, x \mapsto \langle\,\underline{3}\,\rangle)\,\rangle\ \star\ \mathsf{fork}\langle\,\underline{0}\,;\,f(f\,\underline{0})\,;\,\xi\,\rangle\cdot\diamond \\
\succ^* & \langle\,\underline{3}\,\rangle\ \star\ \mathsf{fork}\langle\,\underline{0}\,;\,f(f\,\underline{0})\,;\,\xi\,\rangle\cdot\diamond \\
\succ & \langle\,f(f\,\underline{0})\,;\,\xi\,\rangle\ \star\ \diamond \\
\succ & \langle\,f\,;\,\xi\,\rangle\ \star\ \mathsf{arg}\langle\,f\,\underline{0}\,;\,\xi\,\rangle\cdot\diamond \\
\succ & \langle\,\lambda y.\mathrm{add}\,x\,y\,;\,x \mapsto \langle\,\underline{3}\,\rangle\,\rangle\ \star\ \mathsf{arg}\langle\,f\,\underline{0}\,;\,\xi\,\rangle\cdot\diamond \\
\succ & \langle\,f\,\underline{0}\,;\,\xi\,\rangle\ \star\ \mathsf{fun}\langle\,\lambda y.\mathrm{add}\,x\,y\,;\,x \mapsto \langle\,\underline{3}\,\rangle\,\rangle\cdot\diamond \\
\succ^* & \langle\,\underline{3}\,\rangle\ \star\ \mathsf{fun}\langle\,\lambda y.\mathrm{add}\,x\,y\,;\,x \mapsto \langle\,\underline{3}\,\rangle\,\rangle\cdot\diamond \\
\succ^* & \langle\,\underline{6}\,\rangle\ \star\ \diamond
\end{array}
$$

Figure 2: Evaluation of a term in the $\mathsf{CEK_{PCF}}$ abstract machine.

Noticeably, linear dependency allows to get precise information about the functional behaviour of programs without making the language of types too different from the one of simple types (e.g., one does not need to quantify over index variables, as in sized types). The price to pay, however, is that types, and especially higher-order types, need to be *context aware*: when you type u as a subterm of v (see above) you need to know which arguments u will be applied to. Despite this, a genuinely compositional type inference procedure can actually be designed and is the main technical contribution of this paper.

2.1 Linearity, Abstract Machines and the Complexity of Evaluation

Why dependency, but specially *linearity*, are so useful for complexity analysis? Actually, typing a term using linear dependent types requires finding an upper bound to the number of times each value is copied by its environment, called its *potential*. In the term v from the example above, the variable x is used twice, and accordingly one finds $c < 2$ in (4). Potentials of higher-order values occurring in a term are crucial parameters for the complexity of evaluating the term by abstract mechanisms [11]. The following is an hopefully convincing (but necessarily informal) discussion about why this is the case.

Configurations of abstract machines for the λ-calculus (like Friedman and Felleisen's CEK and Krivine's KAM) can be thought of as being decomposable into two distinct parts:

- First of all, there are *duplicable* entities which are either copied entirely or turned into non-duplicable entities. This includes, in particular, terms in so-called environments. Each (higher-order) duplicable entity is a subterm of the term the computation started from.
- There are *non-duplicable* entities that the machine uses to look for the next redex to be fired. Typically, these entities are the current term and (possibly) the stack. The essential feature of non-duplicable entities is the fact that they are progressively "consumed" by the machine during evaluation: the search for the next redex somehow consists in traversing the non-duplicable entities until a redex is found *or* a duplicable entity needs to be turned into a non-duplicable one.

As an example, consider the process of evaluating the PCF term

$$(\lambda f.\ \mathtt{ifz}\ (f\ \underline{0})\ \mathtt{then}\ \underline{0}\ \mathtt{else}\ f(f\ \underline{0}))((\lambda x.\lambda y.\mathrm{add}\,x\,y)\ \underline{3})$$

by an appropriate generalisation of the CEK machine, see Figure 2. Initially, the whole term is non-duplicable.

4

By travelling into it, the machine finds a first redex u; at that point, $\underline{3}$ becomes duplicable. The obtained closure itself becomes part of the environment ξ, and the machine looks into the body of t, ending up in an occurrence of f, which needs to be replaced by a copy of $\xi(f)$. After an *instantiation step*, a new non-duplicable entity $\xi(f)$ indeed appears. Note that, by an easy combinatorial argument, the number of machine steps necessary to reach f is at most (proportional to) the size of the starting term tu, since reaching f requires consuming non-duplicable entities which can only be created through instantiations. After a copy of $\xi(f)$ becomes non-duplicable, some additional "nonduplicable fuel" becomes available, but not too much: $\lambda y.\mathrm{add}\,x\,y$ is after all a subterm of the initial term.

The careful reader should already have guessed the moral of this story: when analysing the time complexity of evaluation, we could limit ourselves to counting how many *instantiation steps* the machine performs (as opposed to counting *all* machine steps). We claim, on the other hand, that the number of instantiation steps *equals* the sum of potentials of all values appearing in the initial term, something that can be easily inferred from the kind of precise linear typing we were talking about at the beginning of this section.

Summing up, once a dependently linear type has been attributed to a term t, the time complexity of evaluating t can be derived somehow for free: not only an expression bounding the number of instantiation steps performed by an abstract machine evaluating t can be derived, but it is part of the underlying type derivation, essentially. As a consequence, reasoning (automatically or not) about it can be done following the structure of the program.

3 Programs and Types, Formally

In this section, we present the formalism of dℓPCF, a *linear dependent type system* for PCF [27]. Two versions exist: dℓPCF$_\mathsf{N}$ and dℓPCF$_\mathsf{V}$, corresponding to *call-by-name* and *call-by-value* evaluation of terms, respectively. The two type systems are differerent, but the underlying idea is basically the same. We give here the details of the CBV version [13], which better corresponds to widespread intuitions about evaluation, but also provide some indications about the CBN setting [12].

3.1 Terms and Indexes

Terms are given by the usual PCF grammar:

$$s, t, u \ := \ x \ \Big|\ \underline{n} \ \Big|\ t\,u \ \Big|\ \lambda x.t \ \Big|\ \mathrm{p}(t) \ \Big|\ \mathrm{s}(t)$$

$$\mathtt{fix}\ x.t \ \Big|\ \mathtt{ifz}\ t\ \mathtt{then}\ u\ \mathtt{else}\ s.$$

A *value* (denoted by v, w etc.) is either a primitive integer $\underline{n}$, or an abstraction $\lambda x.t$, or a fixpoint $\mathtt{fix}\ x.t$. In addition to the usual terms of the λ-calculus, there are a fixpoint construction, primitive natural numbers with predecessor and successor, and conditional branching with a test for zero. For instance, a simple program computing addition is the following:

$$\mathrm{add}\ =\ \ \mathtt{fix}\ x.\ \lambda y.\lambda z.\ \mathtt{ifz}\ y\ \mathtt{then}\ z\ \mathtt{else}\ \mathrm{s}(x\ (\mathrm{p}(y))\ z)\ .$$

3.1.1 Language of Indexes

As explained informally in Section 2, a type in dℓPCF consists in an annotation of a PCF type, where the annotation consists in some *indexes*. The latter are parametrised by a set of index variables $\mathcal{V} = \{a, b, c, \dots\}$ and an untyped signature Θ of *function symbols*, denoted with f, g, h, etc. We assume Θ contains at least the arithmetic symbols $+$, $-$, 0 and 1, and we write $\underline{n}$ for $1 + \cdots + 1$ (n times). Indexes are then constructed by the following grammar:

$$\mathrm{I}, \mathrm{J}, \mathrm{K} \ ::= \ \ a \ \Big|\ \mathrm{f}(\mathrm{I}_1, \dots, \mathrm{I}_n) \ \Big|\ \sum_{a < \mathrm{I}} \mathrm{J} \ \Big|\ \overset{\mathrm{I},\mathrm{J}}{\underset{a}{\bigtriangleup}}\,\mathrm{K}\ ,$$

where n is the arity of f in Θ. Free and bound index variables are defined as usual, taking care that all occurrences of a in J are bound in both $\sum_{a<\mathrm{I}}$ J and $\bigtriangleup_a^{\mathrm{I},\mathrm{K}}$ J. The substitution of a variable a with J in I is written I$\{$J$/a\}$. Given an *equational program* $\mathcal{E}$ attributing a *meaning* in $\mathbb{N}^n \rightharpoonup \mathbb{N}$ to some symbols of arity n (see Section 5), and a *valuation* ρ mapping index variables to $\mathbb{N}$, the *semantics* $[\![\mathrm{I}]\!]_\rho^{\mathcal{E}}$ of an index I is either a natural number or undefined. Let us describe how we interpret the last two constructions, namely *bounded sums* and *forest cardinalities*.

Bounded sums have the usual meaning: $\sum_{a<\mathrm{I}}$ J is simply the sum of all possible values of J with a taking the values from 0 up to I, excluded. Describing the meaning of forest cardinalities, on the other hand, requires some effort. Informally, the index $\bigtriangleup_a^{\mathrm{I},\mathrm{J}}$ K counts the number of nodes in a forest composed of J trees described using K. Each node in the forest is (uniquely) identified by a natural number, starting form I and visiting the tree in pre-order. The index K has the role of describing the number of children of each forest node, e.g. the number of children of the node 0 is K$\{0/a\}$. More formally, the meaning of a forest cardinality is defined by the following two equations:

$$\bigtriangleup_a^{\mathrm{I},0} \mathrm{K} = 0$$

$$\bigtriangleup_a^{\mathrm{I},\mathrm{J}+1} \mathrm{K} = \left(\bigtriangleup_a^{\mathrm{I},\mathrm{J}} \mathrm{K} \right) + 1 + \left(\bigtriangleup_a^{\mathrm{I}+1+\bigtriangleup_a^{\mathrm{I},\mathrm{J}} \mathrm{K}, \mathrm{K}\{\mathrm{I}+\bigtriangleup_a^{\mathrm{I},\mathrm{J}} \mathrm{K}/a\}} \mathrm{K} \right)$$

The first equation says that a forest of 0 trees contains no nodes. The second one tells us that a forest of J + 1 trees contains:
- The nodes in the first J trees;
- plus the nodes in the last tree, which are just one plus the nodes in the immediate subtrees of the root, considered themselves as a forest.

Consider the following forest comprising two trees:

and consider an index K with a free index variable a such that K$\{1/a\}$ = 3; K$\{\underline{n}/a\}$ = 2 for $n \in \{2,8\}$; K$\{\underline{n}/a\}$ = 1 when $n \in \{0,6,9,11\}$; and K$\{\underline{n}/a\}$ = 0 when $n \in \{3,4,7,10,12\}$. That is, K describes the number of children of each node. Then $\bigtriangleup_a^{0,2}$ K = 13 since it takes into account the entire forest; $\bigtriangleup_a^{0,1}$ K = 8 since it takes into account only the leftmost tree; $\bigtriangleup_a^{8,1}$ K = 5 since it takes into account only the second tree of the forest; finally, $\bigtriangleup_a^{2,3}$ K = 6 since it takes into account only the three trees (as a forest) within the dashed rectangle.

One may wonder what is the role of forest cardinalities in the type system. Actually, they play a crucial role in the treatment of recursion, where the unfolding of recursive calls produces a tree-like structure whose size is just the number of times the (recursively defined) function will be used *globally*.

Notice that $[\![\mathrm{I}]\!]_\rho^{\mathcal{E}}$ is undefined whenever the equality between I and any natural number cannot be derived from the underlying equational program. In particular, a forest cardinality may be undefined even if all its subterms are defined: as an example I = $\bigtriangleup_a^{0,1}$ 1 has no value, because the corresponding tree consists of an infinite descending chain and its cardinality is infinite. By the way I is the index term describing the structure of the recursive calls induced by the program `fix` $x.t$.

3.1.2 Semantic Judgements

A *constraint* is an inequality on indexes. A constraint $I \leqslant J$ is *valid* for ρ and $\mathcal{E}$ when both $[\![I]\!]_\rho^{\mathcal{E}}$ and $[\![J]\!]_\rho^{\mathcal{E}}$ are defined, and $[\![I]\!]_\rho^{\mathcal{E}} \leqslant [\![J]\!]_\rho^{\mathcal{E}}$. As usual, we can derive a notion of equality and strict inequality from $\leqslant$. A *semantic judgement* is of the form

$$\phi; \Phi \vDash_{\mathcal{E}} I \leqslant J \ ,$$

where Φ is a set of constraints and ϕ is the set of free index variables in Φ, I and J. These semantic judgements are used as axioms in the typing derivations of $\mathsf{d}\ell\mathsf{PCF}$, and the set of constraints Φ, called the *index context*, contains *mainly* some indications of bounds for the free index variables (such as $a < K$). Such a judgement is valid when, for every valuation $\rho : \phi \to \mathbb{N}$, if all constraints in Φ are valid for $\mathcal{E}$ and ρ then so is $I \leqslant J$.

3.2 Types

Remember that $\mathsf{d}\ell\mathsf{PCF}$ is aimed at controlling the complexity of programs. The time complexity *of the evaluation* is thus analysed statically, while typing the term at hand. The grammar for types distinguishes the subclass of *linear types*, which correspond to *non-duplicable* terms (see Section 2.1), and the one of *modal types*, for *duplicable* terms. In $\mathsf{d}\ell\mathsf{PCF}_V$, they are defined as follows:

$$\begin{aligned} A, B &\quad ::= \quad \sigma \multimap \tau; &\text{linear types} \\ \sigma, \tau &\quad ::= \quad [a < I] \cdot A \ \Big| \ \mathtt{Nat}[I, J]. &\text{modal types} \end{aligned}$$

Indeed, CBV evaluation only duplicates values. If such a value has an arrow type then it is a function (either an abstraction or a fixpoint) that can potentially increase the complexity of the whole program if we duplicate it. Hence we may need a bound on the number of times we instantiate it in order to control the complexity. This bound, call the *potential* of the value, is represented by I in the type $[a < I] \cdot (\sigma \multimap \tau)$ (also written $\sigma \overset{a<I}{\multimap} \tau$). As explained in Section 2, $\mathtt{Nat}[I, J]$ is the type of programs evaluating to a natural number in the closed interval $[I, J]$. The potential for natural numbers is not specified, as they can be freely duplicated along CBV evaluation.

3.2.1 Summing types

Intuitively, the modal type $\sigma \equiv [a < I] \cdot A$ is assigned to terms that can be copied I times, the k^{th} copy being of type $A\{\underline{k} - 1/a\}$. For those readers who are familiar with Linear Logic, σ can be thought of as representing the type $A\{0/a\} \otimes \cdots \otimes A\{I - 1/a\}$.

In the typing rules we are going to define, modal types need to be manipulated in an algebraic way. For this reason, two operations on modal types are required. The first one is a binary operation $\uplus$ on modal types. Suppose that $\sigma \equiv [a < I] \cdot A\{a/c\}$ and that $\tau \equiv [a < J] \cdot A\{I + a/c\}$. In other words, σ consists of the first I instances of A, i.e. $A\{0/c\} \otimes \cdots \otimes A\{I - 1/c\}$ while τ consists of the next J instances of A, i.e. $A\{I + 0/c\} \otimes \cdots \otimes A\{I + J - 1/c\}$. Their *sum* $\sigma \uplus \tau$ is naturally defined as a modal type consisting of the first $I + J$ instances of A, i.e. $[c < I + J] \cdot A$. Furthermore, $\mathtt{Nat}[I, J] \uplus \mathtt{Nat}[I, J]$ is just $\mathtt{Nat}[I, J]$. A bounded sum operator on modal types can be defined by generalising the idea above: suppose that

$$\sigma = [b < J] \cdot A \left\{ b + \sum_{d < a} J\{d/a\}/c \right\}.$$

Then its *bounded sum* $\sum_{a < I} \sigma$ is just $[c < \sum_{a < I} J] \cdot A$. Finally, $\sum_{a < I} \mathtt{Nat}[J, K] = \mathtt{Nat}[J, K]$, provided a does not occur free in J nor in K.

3.2.2 Subtyping

Central to $\mathsf{d}\ell\mathsf{PCF}$ is the notion of subtyping. An inequality relation $\sqsubseteq$ between (linear or modal) types can be defined using the formal system in Fig. 3. This relation corresponds to lifting index inequalities to the

$$\dfrac{\begin{array}{c}\phi;\Phi \vDash_{\mathcal{E}} \mathrm{K} \leqslant \mathrm{I}\\ \phi;\Phi \vDash_{\mathcal{E}} \mathrm{J} \leqslant \mathrm{H}\end{array}}{\phi;\Phi \vdash_{\mathcal{E}} \mathtt{Nat}[\mathrm{I},\mathrm{J}] \sqsubseteq \mathtt{Nat}[\mathrm{K},\mathrm{H}]} \qquad \dfrac{\begin{array}{c}\phi;\Phi \vdash_{\mathcal{E}} \varrho \sqsubseteq \sigma\\ \phi;\Phi \vdash_{\mathcal{E}} \tau \sqsubseteq \varphi\end{array}}{\phi;\Phi \vdash_{\mathcal{E}} \sigma \multimap \tau \sqsubseteq \varrho \multimap \varphi}$$

$$\dfrac{\begin{array}{c}(a,\phi);(a < \mathrm{J},\Phi) \vdash_{\mathcal{E}} A \sqsubseteq B\\ \phi;\Phi \vDash_{\mathcal{E}} \mathrm{J} \leqslant \mathrm{I}\end{array}}{\phi;\Phi \vdash_{\mathcal{E}} [a < \mathrm{I}] \cdot A \sqsubseteq [a < \mathrm{J}] \cdot B}$$

Figure 3: Subtyping derivation rules of dℓPCF$_V$.

$$\dfrac{}{\phi;\Phi;\Gamma, x : \sigma \vdash_0^{\mathcal{E}} x : \sigma}(Ax) \qquad \dfrac{\phi;\Phi;\Gamma \vdash_{\mathrm{I}}^{\mathcal{E}} t : \sigma \qquad \phi;\Phi \vdash_{\mathcal{E}} \Delta \sqsubseteq \Gamma \qquad \phi;\Phi \vdash_{\mathcal{E}} \sigma \sqsubseteq \tau \qquad \phi;\Phi \vDash_{\mathcal{E}} \mathrm{I} \leqslant \mathrm{J}}{\phi;\Phi;\Delta \vdash_{\mathrm{J}}^{\mathcal{E}} t : \tau}(Subs)$$

$$\dfrac{(a,\phi);(a < \mathrm{I},\Phi);\Gamma, x : \sigma \vdash_{\mathrm{K}}^{\mathcal{E}} t : \tau}{\phi;\Phi;\sum_{a<\mathrm{I}}\Gamma \vdash_{\mathrm{I}+\sum_{a<\mathrm{I}}\mathrm{K}}^{\mathcal{E}} \lambda x.t : [a < \mathrm{I}] \cdot \sigma \multimap \tau}(\multimap) \qquad \dfrac{\phi;\Phi;\Gamma \vdash_{\mathrm{K}}^{\mathcal{E}} t : [a < 1] \cdot \sigma \multimap \tau \qquad \phi;\Phi;\Delta \vdash_{\mathrm{H}}^{\mathcal{E}} u : \sigma\{0/a\}}{\phi;\Phi;\Gamma \uplus \Delta \vdash_{\mathrm{K}+\mathrm{H}}^{\mathcal{E}} tu : \tau\{0/a\}}(App)$$

$$\dfrac{\phi;\Phi;\Gamma \vdash_{\mathrm{M}}^{\mathcal{E}} t : \mathtt{Nat}[\mathrm{J},\mathrm{K}] \qquad \phi;(\mathrm{J} \leqslant 0,\Phi);\Delta \vdash_{\mathrm{N}}^{\mathcal{E}} u : \tau \qquad \phi;(\mathrm{K} \geqslant 1,\Phi);\Delta \vdash_{\mathrm{N}}^{\mathcal{E}} s : \tau}{\phi;\Phi;\Gamma \uplus \Delta \vdash_{\mathrm{M}+\mathrm{N}}^{\mathcal{E}} \mathtt{ifz}\ t\ \mathtt{then}\ u\ \mathtt{else}\ s : \tau}(If)$$

$$\dfrac{}{\phi;\Phi;\Gamma \vdash_0^{\mathcal{E}} \underline{n} : \mathtt{Nat}[\underline{n},\underline{n}]}(n) \qquad \dfrac{\phi;\Phi;\Gamma \vdash_{\mathrm{M}}^{\mathcal{E}} t : \mathtt{Nat}[\mathrm{I},\mathrm{J}]}{\phi;\Phi;\Gamma \vdash_{\mathrm{M}}^{\mathcal{E}} \mathtt{s}(t) : \mathtt{Nat}[\mathrm{I}+1,\mathrm{J}+1]}(s) \qquad \dfrac{\phi;\Phi;\Gamma \vdash_{\mathrm{M}}^{\mathcal{E}} t : \mathtt{Nat}[\mathrm{I},\mathrm{J}]}{\phi;\Phi;\Gamma \vdash_{\mathrm{M}}^{\mathcal{E}} \mathtt{p}(t) : \mathtt{Nat}[\mathrm{I}-1,\mathrm{J}-1]}(p)$$

$$\dfrac{\begin{array}{c}(b,\phi);(b < \mathrm{H},\Phi);\Gamma, x : [a < \mathrm{I}] \cdot A \vdash_{\mathrm{J}}^{\mathcal{E}} t : [a < 1] \cdot B\\ (a,b,\phi);(a < \mathrm{I}, b < \mathrm{H},\Phi) \vdash_{\mathcal{E}} B\{0/a\}\{\textcircled{\bigcirc}_b^{b+1,a} \mathrm{I} + b + 1/b\} \sqsubseteq A\end{array}}{\phi;\Phi;\sum_{b<\mathrm{H}}\Gamma \vdash_{\mathrm{H}+\sum_{b<\mathrm{H}}\mathrm{J}}^{\mathcal{E}} \mathtt{fix}\ x.t : [a < \mathrm{K}] \cdot B\{0/a\}\{\textcircled{\bigcirc}_b^{0,a} \mathrm{I}/b\}}(Fix)$$

(where $\mathrm{H} = \textcircled{\bigcirc}_b^{0,\mathrm{K}} \mathrm{I}$)

Figure 4: Typing rules of dℓPCF$_V$.

type level. As defined here, $\sqsubseteq$ is a pre-order (i.e. a reflexive and transitive relation), which allows to cope with approximations in the typed analysis of programs. However, in the type inference algorithm we will present in next section only the symmetric closure $\equiv$ of $\sqsubseteq$, called *type equivalence* will be used. This ensures that the type produced by the algorithm is *precise*.

3.2.3 Typing

A typing judgement is of the form

$$\phi;\Phi;\Gamma \vdash_{\mathrm{K}}^{\mathcal{E}} t : \tau,$$

where K is the *weight* of t, that is (informally) the maximal number of substitutions involved in the CBV evaluation of t (including the potential substitutions by t itself in its evaluation context). The index context Φ is as in a semantic judgement (see Section 3.1.2), and Γ is a (term) context assigning a modal type to (at least) each free variable of t. Both sums and bounded sums are naturally extended from modal types to contexts (with, for instance, $\{x : \sigma; y : \tau\} \uplus \{x : \varrho, z : \varphi\} = \{x : \sigma \uplus \varrho; y : \tau; z : \varphi\}$). There might be free index variables in Φ, Γ, τ and K, all of them from ϕ. Typing judgements can be derived from the rules of Figure 4.

Observe that, in the typing rule for the abstraction ($\multimap$), I represents the number of times the value $\lambda x.t$ can be copied. Its weight (that is, the number of substitutions involving $\lambda x.t$ or one of its subterms) is then I plus, for each of these copies, the weight of t. In the typing rule (App), on the other hand, t is used once as a function, without been copied. Its potential needs to be at least 1. The typing rule for the fixpoint is arguably the most complicated one. As a first approximation, assume that only one copy of $\mathtt{fix}\ x.t$ will

8

be used (that is, $K \equiv 1$ a does not occur free in B). To compute the weight of `fix` $x.t$, we need to know the number of times t will be copied during the evaluation, that is the number of nodes in the tree of its recursive calls. This tree is described by the index I (as explained in Section 3.1.1), since each occurrence of x in t stands for a recursive call. It has $H = \bigtriangleup_b^{0,1} I$ nodes. At each node b of this tree, there is a copy of t in which the a^{th} occurrence of x will be replaced by the a^{th} son of b, i.e. by $b + 1 + \bigtriangleup_b^{b+1,a} I$. The types thus have to correspond, which is what the second premise of this rule prescribes. Now if `fix` $x.t$ is in fact aimed at being copied $K \geqslant 0$ times, then all the copies of t are represented in a forest of K trees described by I.

For the sake of simplicity, we present here the type system with an explicit subsumption rule. The latter allows to relax any bound in the types (and the weight), thereby loosing some precision in the information provided by the typing judgement. However, we could alternatively replace this rule by relaxing the premises of all the other ones (which corresponds to the presentation of the type system given in [13], or in [12] for $d\ell PCF_N$). Restricting subtyping to type equivalence amounts to considering types up to index equality in the type system of Figure 4 without the rule $(Subs)$ — this is what we do in the type inference algorithm in Section 4. In this case we say that the typing judgements are *precise*:

Definition 3.1 *A derivable judgement* $\phi; \Phi; \Gamma \vdash_I^{\mathcal{E}} t : \sigma$ *is* precise *if*

$$\phi; \Phi; \Delta \vdash_J^{\mathcal{E}} t : \tau \text{ is derivable} \quad \Longrightarrow \quad \left\{ \begin{array}{l} \phi; \Phi \vdash_{\mathcal{E}} \Delta \sqsubseteq \Gamma \\ \phi; \Phi \vdash_{\mathcal{E}} \sigma \sqsubseteq \tau \\ \phi; \Phi \vDash_{\mathcal{E}} I \leqslant J \end{array} \right.$$

3.2.4 Call-by-value vs. Call-by-name

In $d\ell PCF_N$, the syntax of terms and of indexes is the same as in $d\ell PCF_V$, but the language of types differs:

$$\begin{array}{lll} A, B & ::= & \sigma \multimap A \quad \big| \quad \mathtt{Nat}[I, J]; \quad \text{linear types} \\ \sigma, \tau & ::= & [a < I] \cdot A. \quad \text{modal types} \end{array}$$

Modal types still represent duplicable terms, except that now not only values but any argument to functions can be duplicated. So modal types only occur in negative position in arrow types. In the same way, one can find them in the *context* of any typing judgement,

$$\phi; \Phi; (x_1 : \sigma_1, \ldots, x_n : \sigma_n) \vdash_K^{\mathcal{E}} t : A.$$

When a term is typed, it is *a priori* not duplicable, and its type is linear. It is turned into a duplicable term when it holds the position of the argument in an application. As a consequence, the typing rule (App) becomes the most "expansive" one (for the weight) in $d\ell PCF_N$: the whole context used to type the argument has to be duplicated, whereas in $d\ell PCF_V$ this duplication of context is "anticipated" in the typing rules for values.

The readers who are familiar with linear logic, could have noted that, if we replace modal types by banged types, and we remove all annotations with indexes, then $d\ell PCF_N$ corresponds to the target fragment of the CBN translation from simply-typed λ-calculus to LL, and $d\ell PCF_V$ to the target of the CBV translation [26].

In $d\ell PCF_N$, the weight K of a typing judgement represents the maximal number of substitutions that may occur in the CBN evaluation of t. We do not detail the typing rules of $d\ell PCF_N$ here (they can be found in [12]). However an important remark is that in $d\ell PCF_N$, just like in $d\ell PCF_V$, some semantic judgements can be found in the axioms of a typing derivation, and that every typing rule is reversible (except subsumption). The type inference algorithm for $d\ell PCF_V$ that we present in Section 4 can be easily adapted to $d\ell PCF_N$.

3.3 Abstract Machines

The evaluation of PCF terms can be simulated through an extension KAM_{PCF} of Krivine's abstract machine [24] (for CBN evaluation) or through an extension CEK_{PCF} of the Felleisen and Friedman's CEK machine [15] (for CBV evaluation).

$$
\begin{array}{rclcrcl}
\langle tu\,;\xi\rangle & \star & \pi & > & \langle t\,;\xi\rangle & \star & \mathsf{arg}\langle u\,;\xi\rangle\cdot\pi \\
\langle \lambda x.t\,;\xi\rangle & \star & \mathsf{arg}(\mathsf{c})\cdot\pi & > & \langle t\,;(x\mapsto\mathsf{c})\cdot\xi\rangle & \star & \pi \\
\langle \mathtt{fix}\ x.t\,;\xi\rangle & \star & \pi & > & \langle t\,;(x\mapsto\langle \mathtt{fix}\ x.t\,;\xi\rangle)\cdot\xi\rangle & \star & \pi \\
\langle x\,;\xi\rangle & \star & \pi & > & \xi(x) & \star & \pi \\
\langle \mathtt{ifz}\ t\ \mathtt{then}\ u\ \mathtt{else}\ s\,;\xi\rangle & \star & \pi & > & \langle t\,;\xi\rangle & \star & \mathsf{fork}\langle u\,;s\,;\xi\rangle\cdot\pi \\
\langle \underline{0}\,;\xi'\rangle & \star & \mathsf{fork}\langle t\,;u\,;\xi\rangle\cdot\pi & > & \langle t\,;\xi\rangle & \star & \pi \\
\langle \underline{n{+}1}\,;\xi'\rangle & \star & \mathsf{fork}\langle t\,;u\,;\xi\rangle\cdot\pi & > & \langle u\,;\xi\rangle & \star & \pi \\
\langle \mathsf{s}(t)\,;\xi\rangle & \star & \pi & > & \langle t\,;\xi\rangle & \star & \mathsf{s}\cdot\pi \\
\langle \mathsf{p}(t)\,;\xi\rangle & \star & \pi & > & \langle t\,;\xi\rangle & \star & \mathsf{p}\cdot\pi \\
\langle \underline{n}\,;\xi\rangle & \star & \mathsf{s}\cdot\pi & > & \langle \underline{n{+}1}\rangle & \star & \pi \\
\langle \underline{n}\,;\xi\rangle & \star & \mathsf{p}\cdot\pi & > & \langle \underline{n{-}1}\rangle & \star & \pi \\
\end{array}
$$

Figure 5: $\mathsf{KAM}_{\mathsf{PCF}}$ evaluation rules.

Both these machines have states in the form of *processes*, that are pairs of a *closure* (i.e. a term with an environment defining its free variables) and a *stack*, representing the evaluation context. In the $\mathsf{KAM}_{\mathsf{PCF}}$, these objects are given by the following grammar:

$$
\begin{array}{llcl}
\text{Closures:} & \mathsf{c} & := & \langle t\,;\xi\rangle; \\
\text{Environment:} & \xi & := & \{x_1\mapsto \mathsf{c}_1;\cdots;x_k\mapsto \mathsf{c}_k\}; \\
\text{Stacks:} & \pi & := & \diamond \ \big|\ \mathsf{arg}\langle t\,;\xi\rangle\cdot\pi \ \big|\ \mathsf{s}\cdot\pi \ \big|\ \mathsf{p}\cdot\pi \\
& & & \big|\ \mathsf{fork}\langle t\,;u\,;\xi\rangle\cdot\pi; \\
\text{Processes:} & P & := & \mathsf{c}\star\pi.
\end{array}
$$

When the environment is empty, we may use the notation $\langle t\rangle$ instead of $\langle t\,;\varnothing\rangle$ for closures. The evaluation rules of the $\mathsf{KAM}_{\mathsf{PCF}}$ are given in Figure 5. The fourth evaluation rule is said to be an *instantiation step*: the value of a variable x is replaced by the term x maps to in the underlying environment ξ.

The $\mathsf{CEK}_{\mathsf{PCF}}$ machine, which performs CBV evaluation, is slightly more complex: within closures, the *value closures* are those whose first component is a value: $\mathsf{v}:=\langle v\,;\xi\rangle$ (remember that a *value* v is of the form $\underline{n}$, $\lambda x.t$ or $\mathtt{fix}\ x.t$). Moreover, environments assign only value closures to variables:

$$
\xi := \{x_1\mapsto \mathsf{v}_1;\cdots;x_k\mapsto \mathsf{v}_k\}\ .
$$

The grammar for stacks is the same, with one additional construction $(\mathsf{fun}(\mathsf{v})\cdot\pi)$ that is used to encapsulate a function (lambda abstraction or fixpoint) while its argument is computed. Indeed, the latter cannot be substituted for a variable if it is not a value. The evaluation rules for processes are the same than in Figure 5, except that the second and the third ones are replaced by the following:

$$
\begin{array}{rclcrcl}
\mathsf{v} & \star & \mathsf{arg}(\mathsf{c})\cdot\pi & > & \mathsf{c} & \star & \mathsf{fun}(\mathsf{v})\cdot\pi \\
\mathsf{v} & \star & \mathsf{fun}\langle \lambda x.t\,;\xi\rangle\cdot\pi & > & \langle t\,;x\mapsto \mathsf{v}\cdot\xi\rangle & \star & \pi \\
\mathsf{v} & \star & \mathsf{fun}\langle \mathtt{fix}\ x.t\,;\xi\rangle\cdot\pi & > & \langle t\,;x\mapsto\langle \mathtt{fix}\ x.t\,;\xi\rangle\cdot\xi\rangle & \star & \mathsf{arg}(\mathsf{v})\cdot\pi
\end{array}
$$

An example of the evaluation of a term in the $\mathsf{CEK}_{\mathsf{PCF}}$ machine has been shown in Figure 2.

We say that a term t *evaluates* to u in an abstract machine when $\langle t\rangle\star\diamond > \langle u\,;\xi\rangle\star\diamond$. Observe that if t is a closed term then u is necessarily a value. We write $t\Downarrow_{\mathsf{N}}^{n} u$ whenever the $\mathsf{KAM}_{\mathsf{PCF}}$ evaluates t to u in exactly n steps, and $t\Downarrow_{\mathsf{v}}^{n} u$ when the same holds for the $\mathsf{CEK}_{\mathsf{PCF}}$ (we may also omit the exponent n when the number of steps is not relevant).

Abstract Machines and Weight The *weight* of a typable term was informally presented as the number of instantiation steps in its evaluation. Abstract machines enable a more precise formulation of this idea:

Fact 1 *1. If $t\Downarrow_{\mathsf{N}} u$, and $\vdash_{\mathrm{I}}^{\mathcal{E}} t:A$ is derivable in $\mathsf{d\ell PCF_N}$, then $[\![\mathrm{I}]\!]^{\mathcal{E}}$ is an upper bound for the number of instantiation steps in the evaluation of t by the $\mathsf{KAM}_{\mathsf{PCF}}$.*

2. *If $t \Downarrow_{\mathsf{v}} u$, and $\vdash_{\mathrm{I}}^{\mathcal{E}} t : [a < 1] \cdot A$ is derivable in $\mathsf{d}\ell\mathsf{PCF_V}$, then $\llbracket \mathrm{I} \rrbracket^{\mathcal{E}}$ is an upper bound for the instantiation steps in the evaluation of t in the $\mathsf{CEK_{PCF}}$.*

This can be shown by extending the notion of weight and of typing judgement to stacks and processes [12, 13], and is the main ingredients for proving Intensional Soundness (see Section 3.4).

3.4 Key Properties

In this section we briefly recall the main properties of $\mathsf{d}\ell\mathsf{PCF}$, arguing for its relevance as a methodology for complexity analysis. We give the results for $\mathsf{d}\ell\mathsf{PCF_V}$, but they also hold for $\mathsf{d}\ell\mathsf{PCF_N}$ (all proofs can be found in [12, 13]).

The subject reduction property guaranties as usual that typing is correct with respect to term evaluation, but specifies also that the weight of a term cannot increase along reduction:

Proposition 3.1 (Subject Reduction) *For any PCF-terms t, u, if $\phi; \Phi; \varnothing \vdash_{\mathrm{I}}^{\mathcal{E}} t : \tau$ is derivable in $\mathsf{d}\ell\mathsf{PCF_V}$, and if $t \to u$ in CBV, then $\phi; \Phi; \varnothing \vdash_{\mathrm{J}}^{\mathcal{E}} u : \tau$ is also derivable for some J such that $\phi; \Phi \vDash_{\mathcal{E}} \mathrm{J} \leqslant \mathrm{I}$.*

As a consequence, the weight does not tell us much about the number of reduction steps bringing a (typable) term to its normal form. So-called *Intensional Soundness*, on the other hand, allows to deduce some sensible information about the time complexity of evaluating a typable PCF program *in an abstract machine* from its $\mathsf{d}\ell\mathsf{PCF}$ typing judgement.

Proposition 3.2 (Intensional Soundness) *For any term t, if $\vdash_{\mathrm{K}}^{\mathcal{E}} t : \mathtt{Nat}[\mathrm{I}, \mathrm{J}]$ is derivable in $\mathsf{d}\ell\mathsf{PCF_V}$, then t evaluates to $\underline{\mathbf{n}}$ in k steps in the $\mathsf{CEK_{PCF}}$, with $\llbracket \mathrm{I} \rrbracket^{\mathcal{E}} \leqslant n \leqslant \llbracket \mathrm{J} \rrbracket^{\mathcal{E}}$ and $k \leqslant |t| \cdot (\llbracket \mathrm{K} \rrbracket^{\mathcal{E}} + 1)$.*

Intensional Soundness guarantees that the evaluation of any program typable in $\mathsf{d}\ell\mathsf{PCF}$ takes (at most) a number of steps directly proportional to both its syntactic size and its weight. A similar theorem holds when t has a functional type: if, as an example, the type of t is $\mathtt{Nat}[a] \stackrel{1}{\multimap} \mathtt{Nat}[\mathrm{J}]$, then K is parametric on a and $(|t| + 2) \cdot (\llbracket \mathrm{K} \rrbracket^{\mathcal{E}} + 1)$ is an upper bound on the complexity of evaluating t on any integer a.

But is $\mathsf{d}\ell\mathsf{PCF}$ powerful enough to type natural complexity bounded programs? Actually, it is as powerful as PCF itself, since any PCF type derivation can be turned into a $\mathsf{d}\ell\mathsf{PCF}$ one (for an expressive enough equational program), as formalised by the type inference algorithm (Section 4). We can make this statement even more precise for terms of base or first order type, provided two conditions are satisfied:
- On the one hand, the equational program $\mathcal{E}$ needs to be *universal*, meaning that every partial recursive function is representable by some index term. This can be guaranteed, as an example, by the presence of a universal program in $\mathcal{E}$.
- On the other hand, all *true* statements in the form $\phi; \Phi \vDash_{\mathcal{E}} \mathrm{I} \leqslant \mathrm{J}$ must be "available" in the type system for completeness to hold. In other words, one cannot assume that those judgements are derived in a given (recursively enumerable) formal system, because this would violate Gödel's Incompleteness Theorem. In fact, in $\mathsf{d}\ell\mathsf{PCF}$ completeness theorems are *relative* to an oracle for the truth of those assumptions, which is precisely what happens in Floyd-Hoare logics [9].

Proposition 3.3 (Relative Completeness) *If $\mathcal{E}$ is universal, then for any PCF term t,*
1. *if $t \Downarrow_{\mathsf{v}}^{k} \underline{\mathbf{m}}$, then $\vdash_{\underline{k}}^{\mathcal{E}} t : \mathtt{Nat}[\underline{m}]$ is derivable in $\mathsf{d}\ell\mathsf{PCF_V}$.*
2. *if, for any $n \in \mathbb{N}$, there exist k_n, m_n such that $t \underline{\mathbf{n}} \Downarrow_{\mathsf{v}}^{k_n} \underline{\mathbf{m_n}}$, then there exist I and J such that*

$$a; \varnothing; \varnothing \vdash_{\mathrm{I}}^{\mathcal{E}} t : [b < 1] \cdot (\mathtt{Nat}[a] \multimap \mathtt{Nat}[\mathrm{J}])$$

is derivable in $\mathsf{d}\ell\mathsf{PCF_V}$, with $\vDash_{\mathcal{E}} \mathrm{J}\{\underline{n}/a\} = \underline{m_n}$ and $\vDash_{\mathcal{E}} \mathrm{I}\{\underline{n}/a\} \leqslant \underline{k_n}$ for all $n \in \mathbb{N}$.

The careful reader should have noticed that there is indeed a gap between the lower bound provided by completeness and the upper bound provided by soundness: this is indeed the reason why our complexity analysis is only meaningful in an asymptotic sense. Sometimes, however, programs with the same asymptotic behavior *can* indeed be distinguished, e.g. when their size is small relative to the constants in their weight.

In the next section, we will see how to make a concrete use of Relative Completeness. Indeed, we will describe an algorithm that, given a PCF term, returns a dℓPCF judgement $\vdash^{\mathcal{E}}_{\mathrm{K}} t : \tau$ for this term, where $\mathcal{E}$ is equational program that is not *universal*, but expressive enough to derive the typing judgement. To cope with the "relative" part of the result (i.e., the very strong assumption that every *true* semantic judgement must be available), the algorithm also return a set of *side conditions* that have to be checked. These side conditions are in fact semantic judgements that act as axioms (of instances of the subsumption rule) in the typing derivation.

4 Relative Type Inference

Given on the one hand soundness and relative completeness of dℓPCF, and on the other undecidability of complexity analysis for PCF programs, one may wonder whether looking for a type inference procedure makes sense at all. As stressed in the Introduction, we will not give a type inference algorithm *per se*, but rather reduce type inference to the problem of checking a set of inequalities modulo an equational program (see Figure 1). This is the reason why we can only claim type inference to be algorithmically solvable in a *relative* sense, i.e. assuming the existence of an oracle for proof obligations.

Why is solving relative type inference useful? Suppose you have a program $t : \mathtt{Nat} \Rightarrow \mathtt{Nat}$ and you want to prove that it works in a number of steps bounded by a polynomial $p : \mathbb{N} \to \mathbb{N}$ (e.g., $p(x) = 4 \cdot x + 7$). You could of course proceed by building a dℓPCF type derivation for t by hand, or even reason directly on the complexity of t. Relative type inference simplifies your life: it outputs an equational program $\mathcal{E}$, a *precise* type derivation for t whose conclusion is $a; \varnothing; \varnothing \vdash^{\mathcal{E}}_{\mathrm{I}} t : \mathtt{Nat}[a] \overset{1}{\multimap} \mathtt{Nat}[\mathrm{J}]$ and a set $\mathcal{I}$ of inequalities on the same signature as the one of $\mathcal{E}$. Your original problem, then, is reduced to verifying $\models^{\mathcal{E}} \mathcal{I} \cup \{\mathrm{I} \leqslant p(a)\}$. This is arguably an easier problem than the original one: first of all, it has nothing to do with complexity analysis but is rather a problem about the *value* of arithmetical expressions. Secondly it only deals with first-order expressions.

4.1 An Informal Account

From the brief discussion in Section 2, it should be clear that devising a compositional type inference procedure for dℓPCF is nontrivial: the type one assigns to a subterm heavily depends on the ways the rest of the program uses the subterm. The solution we adopt here consists in allowing the algorithm to return *partially unspecified* equational programs: $\mathcal{E}$ as produced in output by $\mathcal{T}$ gives meaning to all the symbols in the output type derivation *except* those occurring in negative position in its conclusion.

To better understand how the type inference algorithm works, let us consider the following term t:

$$uv = (\lambda x.\lambda y.x(xy))(\lambda z.\mathtt{s}(z)).$$

The subterm u can be given type $(\mathtt{Nat} \Rightarrow \mathtt{Nat}) \Rightarrow \mathtt{Nat} \Rightarrow \mathtt{Nat}$ in PCF, while v has type $\mathtt{Nat} \Rightarrow \mathtt{Nat}$. This means t as a whole has type $\mathtt{Nat} \Rightarrow \mathtt{Nat}$ and computes the function $x \mapsto 2 \cdot x$. The type inference algorithm proceeds by giving types to u and to v separately, then assembling the two into one. Suppose we start with v. The type inference algorithm refines $\mathtt{Nat} \Rightarrow \mathtt{Nat}$ into $\sigma = \mathtt{Nat}[\mathrm{f}(a,b)] \overset{b<\mathrm{h}(a)}{\multimap} \mathtt{Nat}[\mathrm{g}(a,b)]$ and the equational program $\mathcal{A}_v$, which gives meaning to g in terms of f:

$$\mathrm{g}(a,b) = \mathrm{f}(a,b) + 1.$$

Observe how both f and h are not specified in $\mathcal{A}_v$, because they appear in *negative* position in σ: $\mathrm{f}(a,b)$ intuitively corresponds to the argument(s) v will be applied to, while $\mathrm{h}(a)$ is the number of times v will be used. Notice that everything is parametrised on a, which is something like a global parameter that will later

be set as the input to t. The function u, on the other hand, is given type

$$(\text{Nat}[\text{p}(a,b,c)] \overset{c<\text{j}(a,b)}{\multimap} \text{Nat}[\text{q}(a,b,c)])$$
$$\overset{b<\text{k}(a)}{\multimap}$$
$$\text{Nat}[\text{l}(a,b,c)] \overset{c<\text{m}(a,b)}{\multimap} \text{Nat}[\text{n}(a,b,c)].$$

The newly introduced function symbols are subject to the following equations:

$$\text{j}(a,b) = 2 \cdot \text{m}(a,b);$$
$$\text{n}(a,b,c) = \text{q}(a,b,2c);$$
$$\text{p}(a,b,2c) = \text{q}(a,b,2c+1);$$
$$\text{p}(a,b,2c+1) = \text{l}(a,b,2c).$$

Again, notice that some functions are left unspecified, namely l, m, q and k. Now, a type for uv can be found by just combining the types for u and v, somehow following the typing rule for applications. First of all, the number of times u needs to be copied is set to 1 by the equation $\text{k}(a) = 1$. Then, the matching symbols of u and v are defined one in terms of the others:

$$\text{q}(a,0,b) = \text{g}(a,b);$$
$$\text{f}(a,b) = \text{p}(a,0,b);$$
$$\text{h}(a) = \text{j}(a,0).$$

This is the last step of type inference, so it is safe to stipulate that $\text{m}(a,0) = 1$ and that $\text{l}(a,0,c) = a$, thus obtaining a fully specified equational program $\mathcal{E}$ and the following type τ for t:

$$\text{Nat}[a] \overset{c<1}{\multimap} \text{Nat}[\text{n}(a,0,c)].$$

As an exercise, the reader can verify that the equational program above allows to verify that $\text{n}(a,0,0) = a+2$, and that

$$a; \varnothing; \varnothing \vdash_2^{\mathcal{E}} t : \tau.$$

4.2 Preliminaries

Before embarking on the description of the type inference algorithms, some preliminary concepts and ideas need to be introduced, and are the topic of this section.

4.2.1 Getting Rid of Subsumption

The type inference algorithm takes in input a PCF term t, and returns a typing judgement $\mathcal{J}$ for t, together with a set $\mathscr{R}$ of so-called *side conditions*. We will show below that $\mathcal{J}$ is derivable *iff* all the side conditions in $\mathscr{R}$ are valid. Moreover, in this case $\mathcal{J}$ is *precise* (see Definition 3.1): all occurrences of the base type $\text{Nat}[\text{I},\text{J}]$ are in fact of the form $\text{Nat}[\text{I}]$, and the weight and all potentials H occurring in a sub-type $[a < \text{H}] \cdot A$ are kept as low as possible. Concretely, this means that there is a derivation for $\mathcal{J}$ in which the subsumption rule is restricted to the following form:

$$\frac{\phi; \Phi; \Gamma \vdash_{\text{I}}^{\mathcal{E}} t : \sigma \qquad \begin{array}{c} \phi; \Phi \vdash_{\mathcal{E}} \Delta \equiv \Gamma \\ \phi; \Phi \vdash_{\mathcal{E}} \sigma \equiv \tau \\ \phi; \Phi \vDash_{\mathcal{E}} \text{I} = \text{J} \end{array}}{\phi; \Phi; \Delta \vdash_{\text{J}}^{\mathcal{E}} t : \tau}$$

The three premises on the right boil down to a set of semantic judgements of the form $\{\phi; \Phi \vDash_{\mathcal{E}} \text{K}_i = \text{H}_i\}$ (see Figure 3), where the K_i's are indexes occurring in σ or Δ (or I itself) and the H_i's occur in τ or Γ (or are J itself). If the equalities $\text{K}_i = \text{H}_i$ can all be derived from $\mathcal{E}$, then the three premises on the right are equivalent to the conjunction (on i) of the following properties:

$$\text{``}\llbracket H_i \rrbracket_\rho^{\mathcal{E}} \text{ is defined for any } \rho : \phi \to \mathbb{N} \text{ satisfying } \Phi\text{''}$$

(see Section 3.1.2). Given $\mathcal{E}$, this property (called a *side condition*) is denoted by $\phi; \Phi \vDash H_i \downarrow$. Actually the type inference algorithm does *not* verify any semantic or subtyping judgement coming from (instances of) the subsumption rule. Instead, it turns all index equivalences $H_i = K_i$ into rewriting rules in $\mathcal{E}$, and put all side conditions $\phi; \Phi \vDash H_i \downarrow$ in $\mathscr{R}$. If every side condition in $\mathscr{R}$ is true for $\mathcal{E}$, we write $\mathcal{E} \Vdash \bigwedge \mathscr{R}$. Informally, this means that all subsumptions assumed by the algorithm are indeed valid.

4.2.2 Function Symbols

Types and judgements manipulated by our type inference algorithm have a very peculiar shape. In particular, not every index term is allowed to appear in types, and this property will be crucial when showing soundness and completeness of the algorithm itself:

Definition 4.1 (Primitive Types) *A type is* primitive for ϕ *when it is on the form* $\mathtt{Nat}[f(\phi)]$, *or* $A \multimap B$ *with* A *and* B *primitive for* ϕ, *or* $[a < f(\phi)] \cdot A$ *with* $a \notin \phi$ *and* A *primitive for* $a; \phi$. *A type is said to be* primitive *when it is primitive for some* ϕ.

As an example, a primitive type for $\phi = a; b$ is $\mathtt{Nat}[f(a,b,c)] \overset{c < g(a,b)}{\multimap} \mathtt{Nat}[h(a,b,c)]$. Informally, then, a type is primitive when the only allowed index terms are function symbols (with the appropriate arity).

4.2.3 Equational Programs

The equational program our algorithm constructs is in fact a *rewriting program*: every equality corresponds to the (partial) definition of a function symbol, and we may write it $f(a_1, \ldots, a_k) := J$ (where all free variables of J are in $\{a_1, \ldots, a_k\}$). If there is no such equation in the rewriting program, we say that f is *unspecified*.

An equational program $\mathcal{E}$ is *completely specified* if it allows to deduce a precise semantic (namely a partial recursive function) for each symbol of its underlying signature (written $\Sigma_{\mathcal{E}}$), i.e. none of the symbols in $\Sigma_{\mathcal{E}}$ are unspecified. In other words: a completely specified equational programs has only one *model*. On the other hand, a *partially specified* equational program (i.e. a program where symbols can possibly be unspecified) can have many models, because partial recursive functions can be assigned to function symbols in many different ways, all of them consistent with its equations. Up to now, we only worked with completely specified programs, but allowing the possibility to have unspecified symbols is crucial for being able to describe the type inference algorithm in a simple way. In the following, $\mathcal{E}$ and $\mathcal{F}$ denote completely specified equational programs, while $\mathcal{A}$ and $\mathcal{B}$ denote rewriting programs that are only partially specified.

Definition 4.2 (Model of a Rewriting Program) *An* interpretation μ *of* $\mathcal{A}$ *in* $\mathcal{E}$ *is simply a map from unspecified symbols of* $\mathcal{A}$ *to indexes on the signature* $\Sigma_{\mathcal{E}}$, *such that if* f *has arity* n, *then* $\mu(f)$ *is a term in* $\Sigma_{\mathcal{E}}$ *with free variables from* $\{a_1, \ldots, a_n\}$. *When such an interpretation is defined, we say that* $\mathcal{E}$ *is a* model *of* $\mathcal{A}$, *and we write* $\mu : \mathcal{E} \vDash \mathcal{A}$.

Notice that such an interpretation can naturally be extended to arbitrary index terms on the signature $\Sigma_{\mathcal{A}}$, and we assume in the following that a rewriting program and its model have disjoint signatures.

Definition 4.3 (Validity in a Model) *Given* $\mu : \mathcal{E} \vDash \mathcal{A}$, *we say that a semantic judgement* $\phi; \Phi \vDash_{\mathcal{A}} I \leqslant J$ *is* valid in the model *(notation:* $\phi; \Phi \vDash_\mu I \leqslant J$*) when* $\phi; \Phi \vDash_{\mathcal{F}} I \leqslant J$ *where* $\mathcal{F} = \mathcal{A} \cup \mathcal{E} \cup \{f := \mu(f) \mid f$ *is unspecified in* $\mathcal{A}\}$. *This definition is naturally extended to side conditions (with* $\mu \Vdash \bigwedge \mathscr{R}$ *standing for* $\mathcal{F} \Vdash \bigwedge \mathscr{R}$*).*

Please note that if $\mathcal{A}$ is a completely specified rewriting program, then any model $\mu : \mathcal{E} \vDash \mathcal{A}$ has an interpretation μ with an empty domain, and $\mu \Vdash \bigwedge \mathscr{R}$ iff $\mathcal{A} \Vdash \bigwedge \mathscr{R}$ (still assuming that $\Sigma_{\mathcal{A}}$ and $\Sigma_{\mathcal{E}}$ are disjoint). Now, suppose that $\mu : \mathcal{E} \vDash \mathcal{A}$ and that the index terms I (from $\Sigma_{\mathcal{E}}$) and J (from $\Sigma_{\mathcal{A}}$) are such that $\phi; \Phi \vDash_{\mathcal{E}} I \leqslant \mu(J)$. Then, with a slight abuse of notation, we simply write $\phi; \Phi \vDash_\mu I \leqslant J$. The same notation can be extended to types and judgements.

As already mentioned, the equational programs handled by our type inference algorithm are not necessarily completely specified. Function symbols which are not specified are precisely those occurring in "negative position" in the judgement produced in output. This invariant will be very useful and is captured by the following definition:

Definition 4.4 (Positive and Negative Symbols) *Given a primitive type τ, the sets of its* positive *and* negative *symbols (denoted by τ^+ and τ^- respectively) are defined inductively by*

$$\begin{array}{rclcrcl}
\mathtt{Nat}[\mathrm{i}(\phi)]^+ & = & \{\mathrm{i}\}; & & \mathtt{Nat}[\mathrm{i}(\phi)]^- & = & \varnothing; \\
[a < \mathrm{h}(\phi)] \cdot (\sigma \multimap \tau)^+; & = & \sigma^- \cup \tau^+; & & [a < \mathrm{h}(\phi)] \cdot (\sigma \multimap \tau)^-; & = & \{\mathrm{h}\} \cup \sigma^+ \cup \tau^-.
\end{array}$$

Then the set of positive *(resp. negative) symbols of a judgement $\phi; \Phi; (x_i : \sigma_i)_{i \leqslant n} \vdash^{\mathcal{E}}_{\mathrm{I}} t : \tau$ is the union of all negative (resp. positive) symbols of the σ_i's and all positive (resp. negative) symbols of τ.*

Polarities in $\{+, -\}$ are indicated with symbols like p, q. Given such a p, the opposite polarity is $\neg p$.

Definition 4.5 (Specified Symbols, Types and Judgments) *Given a set of function symbols $\mathcal{S}$, a symbol f is said to be $(\mathcal{S}, \mathcal{A})$-specified when there is a rule $\mathrm{f}(\phi) := \mathrm{J}$ in $\mathcal{A}$ such that any function symbol appearing in J is either f itself, or in $\mathcal{S}$, or a symbol that is $(\mathcal{S} \cup \{\mathrm{f}\}, \mathcal{A}\backslash\{\mathrm{f}(\phi) := \mathrm{J}\})$-specified. Remember that when there is no rule $\mathrm{f}(\phi) := \mathrm{J}$ in $\mathcal{A}$ the symbol f is* unspecified *in $\mathcal{A}$. A primitive type σ is said to be $(p, \mathcal{S}, \mathcal{A})$-specified when all function symbols in σ^p are $(\mathcal{S}, \mathcal{A})$-specified and all symbols in $\sigma^{\neg p}$ are unspecified. A judgement $\phi; \Phi; \Gamma \vdash^{\mathcal{A}}_{\mathrm{I}} t : \tau$ is correctly specified when τ and all types in Γ are primitive for ϕ, and τ is $(+, \mathcal{N}, \mathcal{A})$-specified, and all types in Γ are $(-, \mathcal{N}, \mathcal{A})$-specified, and all function symbols in I are $(\mathcal{N}, \mathcal{A})$-specified where $\mathcal{N}$ is the set of negative symbols of the judgement.*

In other words, a judgement is correctly specified if the underlying equational program (possibly recursively) defines all symbols in positive position depending on those in negative position.

4.3 The Structure of the Algorithm

The type inference algorithm that we develop here receives in input a PCF term t and returns a dℓPCF judgement $\varnothing \vdash^{\mathcal{E}}_{\mathrm{K}} t : \tau$ for it, together with a set of side conditions $\mathscr{R}$. We will prove that it is *correct*, in the sense that the typing judgement is derivable *iff* the side conditions hold. The algorithm proceeds as follows:

1. Compute d_{PCF}, a PCF type derivation for t;
2. Proceeding by structural induction on d_{PCF}, construct a dℓPCF derivation for t (call it d_{v}) and the corresponding set of side conditions $\mathscr{R}$;
3. Returns $\mathscr{R}$ and the conclusion of d_{v}.

The *skeleton* $([\sigma])$ (or $([A])$) of a modal type σ (*resp.* of a linear type A) is obtained by erasing all its indexes (and its bounds $[a < \mathrm{I}]$). The *skeleton* of a dℓPCF derivation is obtained by replacing each type by its skeleton, and erasing all the subsumption rules. In PCF the type inference problem is decidable, and the Step 1. raises no difficulty: actually, one could even assume that the type d_{PCF} attributes to t is principal. The core of the algorithm is of course Step 2. In Section. 4.5 we will define a recursive algorithm GEN that build d_{v} and $\mathscr{R}$ by annotating d_{PCF}. The algorithm GEN itself relies on some auxiliary algorithms, which will be described in Section 4.4 below.

All auxiliary algorithms we will talk about have the ability to generate fresh variables and function symbols. Strictly speaking, then, they should take a counter (or anything similar) as a parameter, but we elide this for the sake of simplicity. Also we consider that we assume the existence of a function $\alpha(\phi; T)$ that, given a set of index variables ϕ and a PCF type T, returns a modal type τ primitive for ϕ, containing only fresh function symbols, and such that $([\tau]) = T$.

4.4 Auxiliary Algorithms and Linear Logic

The design of systems of linear dependent types such as dℓPCF$_{\mathsf{V}}$ and dℓPCF$_{\mathsf{N}}$ is strongly inspired by BLL, itself a restriction of linear logic. Actually, the best way to present the type inference algorithm consists in

first of all introducing four auxiliary algorithms, each corresponding to a principle regulating the behaviour of the exponential connectives in linear logic. Notice that these auxiliary algorithms are the main ingredients of both $\mathsf{d\ell PCF_V}$ and $\mathsf{d\ell PCF_N}$ type inference. Consistently to what we have done so far, we will prove and explain them with $\mathsf{d\ell PCF_V}$ in mind. All the auxiliary algorithm we will talk about in this section will take a tuple of $\mathsf{d\ell PCF_V}$ types as first argument; we assume that all of them have the same skeleton and, moreover, that all index terms appearing in them are pairwise distinct.

Dereliction. Dereliction is the following principle: any duplicable object (say, of type $!A$) can be made linear (of type A), that is to say $!A \rightarrow A$. In $\mathsf{d\ell PCF}$, being duplicable means having a modal type, which also contains some quantitative information, namely how many times the object can be duplicated, at most. In $\mathsf{d\ell PCF}$, dereliction can be simply seen as the principle $[a < 1] \cdot A \rightarrow A\{0/a\}$, and is implicitly used in the rules (App) and (Fix). Along the type inference process, as a consequence, we often need to create "fresh instances" of dereliction in the form of pairs of types being in the correct semantic relation. This is indeed possible:

Lemma 4.1 *There is an algorithm* DER *such that given two types* τ *(primitive for* ϕ*) and* σ *(primitive for* a, ϕ*) of the same skeleton,* $\mathrm{DER}((\sigma, \tau); a; \phi; \Phi; p) = (\mathcal{A}, \mathcal{R})$ *where:*
1. *for every* $\mathcal{E} \supseteq \mathcal{A}$*, if* $\mathcal{E} \Vdash \bigwedge \mathcal{R}$ *then* $\phi; \Phi \vdash_{\mathcal{E}} \sigma\{0/a\} \equiv \tau$*;*
2. *whenever* $\phi; \Phi \vdash_{\mathcal{E}} \varrho\{0/a\} \equiv \xi$ *where* $(\![\varrho]\!) \equiv (\![\sigma]\!)$*, there is* $\mu : \mathcal{E} \models \mathcal{A}$ *such that* $\phi; \Phi \vdash_{\mu} \sigma \equiv \varrho$*,* $\phi; \Phi \vdash_{\mu} \tau \equiv \xi$*, and* $\mu \Vdash \bigwedge \mathcal{R}$*;*
3. σ *is* $(p, \tau^p, \mathcal{A})$*-specified and* τ *is* $(\neg p, \sigma^{\neg p}, \mathcal{A})$*-specified.*

Proof. The algorithm DER is defined by recursion on the structure of σ:
- Here are two base cases:

$$\mathrm{DER}((\mathtt{Nat}[\mathrm{f}(a, \phi)], \mathtt{Nat}[\mathrm{g}(\phi)]); a; \phi; \Phi; +) = (\{\mathrm{g}(\phi) := \mathrm{f}(0, \phi)\}, \{\phi, \Phi \models \mathrm{f}(0, \phi) \downarrow\})$$
$$\mathrm{DER}((\mathtt{Nat}[\mathrm{f}(a, \phi)], \mathtt{Nat}[\mathrm{g}(\phi)]); a; \phi; \Phi; -) = (\{\mathrm{f}(a, \phi) := \mathrm{g}(\phi)\}, \{\phi, \Phi \models \mathrm{g}(\phi) \downarrow\})$$

- Inductive cases are slightly more difficult:

$$\mathrm{DER}((\sigma_1 \overset{b < \mathrm{f}(a, \phi)}{\multimap} \sigma_2, \tau_1 \overset{b < \mathrm{g}(\phi)}{\multimap} \tau_2); a; \phi; \Phi; +) = (\mathcal{A}_1^- \cup \mathcal{A}_2^+ \cup \{\mathrm{g}(\phi) := \mathrm{f}(0, \phi)\},$$
$$\mathcal{R}_1^- \cup \mathcal{R}_2^+ \cup \{\phi, \Phi \models \mathrm{f}(0, \phi) \downarrow\});$$

$$\mathrm{DER}((\sigma_1 \overset{b < \mathrm{f}(a, \phi)}{\multimap} \sigma_2, \tau_1 \overset{b < \mathrm{g}(\phi)}{\multimap} \tau_2); a; \phi; \Phi; -) = (\mathcal{A}_1^+ \cup \mathcal{A}_2^- \cup \{\mathrm{f}(a, \phi) := \mathrm{g}(\phi)\},$$
$$\mathcal{R}_1^- \cup \mathcal{R}_2^+ \cup \{\phi, \Phi \models \mathrm{g}(\phi) \downarrow\}).$$

where:

$$\mathrm{DER}((\sigma_1, \tau_1); a; b, \phi; \Phi, b < \mathrm{f}(0, \phi); +) = (\mathcal{A}_1^-, \mathcal{R}_1^-);$$
$$\mathrm{DER}((\sigma_1, \tau_1); a; b, \phi; \Phi, b < \mathrm{g}(\phi); -) = (\mathcal{A}_1^+, \mathcal{R}_1^+);$$
$$\mathrm{DER}((\sigma_2, \tau_2); a; b, \phi; \Phi, b < \mathrm{f}(0, \phi); +) = (\mathcal{A}_2^+, \mathcal{R}_2^+);$$
$$\mathrm{DER}((\sigma_2, \tau_2); a; b, \phi; \Phi, b < \mathrm{g}(\phi); -) = (\mathcal{A}_2^-, \mathcal{R}_2^-).$$

Let us now prove the Lemma by induction on the structure of σ:
- If $\sigma = \mathtt{Nat}[\mathrm{f}(0, \phi]$, then the thesis can be easily reached;
- Suppose that $\sigma = \sigma_1 \overset{b < \mathrm{f}(a, \phi)}{\multimap} \sigma_2$ and that $\tau = \tau_1 \overset{b < \mathrm{g}(\phi)}{\multimap} \tau_2$. For simplicity, suppose that $p = +$ (the case $p = -$ is analogous):
  1. If $\mathcal{E} \supseteq \mathcal{A}$, then $\mathcal{E} \supseteq \mathcal{A}_1^-$ and $\mathcal{E} \supseteq \mathcal{A}_2^+$. Moreover, if $\mathcal{E} \Vdash \bigwedge \mathcal{R}$ then, of course $\mathcal{E} \Vdash \bigwedge \mathcal{R}_1^-$ and $\mathcal{E} \Vdash \bigwedge \mathcal{R}_2^+$. As a consequence, by induction hypothesis, we have

$$\phi, b; \Phi, b < \mathrm{g}(\phi) \vdash_{\mathcal{E}} \sigma_1\{0/a\} \equiv \tau_1;$$
$$\phi, b; \Phi, b < \mathrm{f}(a, \phi) \vdash_{\mathcal{E}} \sigma_2\{0/a\} \equiv \tau_2.$$

16

From $\phi; \Phi \models_{\mathcal{E}} f(0, \phi) \downarrow$ it follows that $\phi; \Phi \models_{\mathcal{E}} g(\phi) = f(0, \phi)$. As a consequence,

$$\phi; \Phi \vdash_{\mathcal{E}} \sigma \equiv \tau.$$

2. If $\phi; \Phi \vdash_{\mathcal{E}} \varrho\{0/a\} \equiv \xi$ where $([\varrho]) \equiv ([\sigma])$, we can safely assume that

$$\varrho = \varrho_1 \overset{b<I}{\multimap} \varrho_2$$
$$\xi = \xi_1 \overset{b<J}{\multimap} \xi_2$$

and moreover, that

$$\phi, b; \Phi, b < I\{0/a\} \vdash_{\mathcal{E}} \varrho_1\{0/a\} \equiv \xi_1\{0/a\};$$
$$\phi, b; \Phi, b < J \vdash_{\mathcal{E}} \varrho_2 \equiv \xi_2;$$
$$\phi; \Phi \models_{\mathcal{E}} I\{0/a\} = J.$$

Assume for the sake of simplicity that $p = +$. Then, from the two judgements above, one obtains:

$$\phi, b; \Phi, b < f(0, \phi) \vdash_{\mathcal{E} \cup \{f(a,\phi):=I\}} \varrho_1\{0/a\} \equiv \xi_1\{0/a\}$$
$$\phi, b; \Phi, b < g(\phi) \vdash_{\mathcal{E} \cup \{g(\phi):=J\}} \varrho_2 \equiv \xi_2$$

By induction hypothesis, there are $\mu_1 : (\mathcal{E} \cup \{f(a, \phi) := I\}) \models \mathcal{A}_1^-$ and $\mu_2 : (\mathcal{E} \cup \{g(\phi) := J\}) \models \mathcal{A}_2^+$ such that

$$\phi, b; \Phi, b < f(0, \phi) \vdash_{\mathcal{E} \cup \{f(a,\phi):=I\}} \sigma_1\{0/a\} \equiv \varrho_1\{0/a\};$$
$$\phi, b; \Phi, b < f(0, \phi) \vdash_{\mathcal{E} \cup \{f(a,\phi):=I\}} \tau_1\{0/a\} \equiv \xi_1\{0/a\};$$
$$\phi, b; \Phi, b < g(\phi) \vdash_{\mathcal{E} \cup \{g(\phi):=J\}} \sigma_2 \equiv \varrho_2;$$
$$\phi, b; \Phi, b < g(\phi) \vdash_{\mathcal{E} \cup \{g(\phi):=J\}} \tau_2 \equiv \xi_2.$$

and $\mu_1 \Vdash \bigwedge \mathscr{R}_1^-$, and $\mu_2 \Vdash \bigwedge \mathscr{R}_2^+$. By way of the push-out technique, we get $\epsilon_1 : \mathcal{E} \models \mathcal{A}_1^-$ and $\epsilon_2 : \mathcal{E} \models \mathcal{A}_2^+$ such that $\epsilon_1 \Vdash \bigwedge \mathscr{R}_1^-$, and $\epsilon_2 \Vdash \bigwedge \mathscr{R}_2^+$. Now, μ is defined as $\mu_1 \cup \mu_2$, plus the assignment of I (after some α-renaming) to f. First of all, notice that $\mu \Vdash \bigwedge \mathscr{R}$, simply because $\phi; \Phi \models_{\mathcal{E}} I\{0/a\} \downarrow$. Moreover, observe that, by easy manipulations of the derivations above, one gets:

$$\phi, b; \Phi, b < I\{0/a\} \vdash_{\mu} \sigma_1\{0/a\} \equiv \varrho_1\{0/a\};$$
$$\phi, b; \Phi, b < I\{0/a\} \vdash_{\mu} \tau_1\{0/a\} \equiv \xi_1\{0/a\};$$
$$\phi, b; \Phi, b < J \vdash_{\mu} \sigma_2 \equiv \varrho_2;$$
$$\phi, b; \Phi, b < J \vdash_{\mu} \tau_2 \equiv \xi_2.$$

3. This is an easy consequence of how σ and τ are built.

This concludes the proof. $\qquad\square$

The algorithm DER works by recursion on the PCF type $([\sigma])$ and has thus linear complexity in $|([\sigma])|$.

Contraction. Another key principle in linear logic is contraction, according to which two copies of a duplicable object can actually be produced, $!A \to !A \otimes !A$. Contraction is used in binary rules like (App) or (If), in the form of the operator $\uplus$. This time, we need an algorithm CTR which takes *three* linear types A, B and C (all of them primitive for (a, ϕ)) and turn them into an equational program and a set of side conditions:

$$\text{CTR}((A, B, C); (I, J); a; \phi; \Phi; p) = (\mathcal{A}, \mathscr{R}).$$

The parameters I and J are index terms capturing the number of times B and C can be copied. A Lemma akin to 4.1 can indeed be proved about CTR.

Lemma 4.2 *There is an algorithm* CTR *such that whenever B and C are two linear types of the same skeleton and primitive for (a, ϕ),* $\mathrm{CTR}((B, C, \mathrm{I}); (\mathrm{J}, a); \phi; \Phi; p; =)(A, \mathcal{A}, \mathscr{R})$, *where:*
- *A is primitive for a, ϕ, and $(\![A]\!) = (\![B]\!) = (\![C]\!)$;*
- *For any $\mathcal{E} \supseteq \mathcal{A}$, $\quad \mathcal{E} \Vdash \bigwedge \mathscr{R}$ iff*

$$\phi; \Phi \vdash_{\mathcal{E}} [a < \mathrm{I} + \mathrm{J}] \cdot A \equiv ([a < \mathrm{I}] \cdot B') \uplus ([a < \mathrm{J}] \cdot C')$$

  for some B' and C' such that $\phi; \Phi, a < \mathrm{I} \vdash_{\mathcal{E}} B' \equiv B$ and $\phi; \Phi, a < \mathrm{J} \vdash_{\mathcal{E}} C' \equiv C$.
- *A is $(p, B^p \cup C^p, \mathcal{A})$-specified, and B and C are $(\neg p, A^{\neg p}, \mathcal{A})$-specified.*

Digging. In linear logic, any duplicable object having type $!A$ can be turned into an object of type $!!A$, namely an object which is the duplicable version of a duplicable object. Digging is the principle according to which this transformation is possible, namely $!A \to !!A$. At the quantitative level, this corresponds to splitting a bounded sum into its summands. This is used in the typing rules for functions, $(\multimap)$ and (Fix).

The auxiliary algorithm corresponding to the digging principle takes two linear types and builds, as usual, a rewriting program and a set of side conditions capturing the fact that the first of the two types is the bounded sum of the second:

$$\mathrm{DIG}((A, B); (\mathrm{I}, \mathrm{J}); \phi; (a, b); \Phi; p) \;=\; (\mathcal{A}, \mathscr{R}).$$

The correctness of DIG can again be proved similarly to what we did in Lemma 4.1, the key statement being that for every $\mathcal{E} \supseteq \mathcal{A}$ such that $\mathcal{E} \Vdash \bigwedge \mathscr{R}$, the following must hold

$$\phi; \Phi \vdash_{\mathcal{E}} [b < \sum_{a < \mathrm{I}} \mathrm{J}] \cdot A \equiv \sum_{a < \mathrm{I}} [b < \mathrm{J}] \cdot C$$

for some C such that $\phi; \Phi, a < \mathrm{I}, b < \mathrm{J} \vdash_{\mathcal{E}} C \equiv B$.

Lemma 4.3 *There is an algorithm* DIG *such that for any linear type B primitive for (a, b, ϕ),* $\mathrm{DIG}((B, \mathrm{I}); (\mathrm{J}, \phi); a; (b, \Phi); p;) = (A, \mathcal{A}, \mathscr{R})$, *where:*
- *A is primitive for b, ϕ, and $(\![A]\!) = (\![B]\!)$;*
- *for any $\mathcal{E} \supseteq \mathcal{A}$, $\quad \mathcal{E} \Vdash \bigwedge \mathscr{R}$ iff*

$$\phi; \Phi \vdash_{\mathcal{E}} [b < \sum_{a < \mathrm{I}} \mathrm{J}] \cdot A \equiv \sum_{a < \mathrm{I}} [b < \mathrm{J}] \cdot B'$$

  for some B' such that $\phi; \Phi, a < \mathrm{I}, b < \mathrm{J} \vdash_{\mathcal{E}} B' \equiv B$;
- *A is $(p, B^p, \mathcal{A})$-specified and B is $(\neg p, A^{\neg p}, \mathcal{A})$-specified.*

Weakening. Weakening means that duplicable objects can also be erased, even when the underlying index is 0. Weakening is useful in the rules (Ax) and (n). Once a fresh $\mathsf{d\ell PCF_V}$ type A is produced, the only thing we need to do is to produce an equational program $\mathcal{A}$ specifying (in an arbitrary way) the symbols in A^p, this way preserving the crucial invariants about the equational programs manipulated by the algorithm. Formally, it means that there is an algorithm WEAK such that

$$\mathrm{WEAK}(A; \phi; a; p) = \mathcal{A},$$

where A is $(p, \varnothing, \mathcal{A})$-specified

Lemma 4.4 *There is an algorithm* WEAK *such that* $\mathrm{WEAK}(T; \phi; a; p) = (A, \mathcal{A})$ *where,*
- *A is primitive for $a; \phi$, and $(\![A]\!) = T$;*
- *for every $\mathcal{E} \supseteq \mathcal{A}$, it holds that $\phi; \Phi, a < 0 \vdash_{\mathcal{E}} A \equiv A$;*
- *A is $(p, \varnothing, \mathcal{A})$-specified.*

Observe how no sets of constraints is produced in output by WEAK, contrarily to DER and CTR.

Some Extra Functions. There are also some extra auxiliary algorithms that are used by the function GEN. Although they are not related to some exponential rule of LL, they are defined in a similar way than the previous one.

The following lemma is used for type checking the (Ax) rule:

Lemma 4.5 *There is an algorithm* EQ *such that, given two types σ and τ primitive for ϕ and of same skeleton,* $\mathrm{EQ}((\tau;\sigma);\phi;\Phi;p) = (\mathcal{A};\mathscr{R})$ *where,*
- *for every $\mathcal{E} \supseteq \mathcal{A}$, $\mathcal{E} \Vdash \bigwedge \mathscr{R} \iff \phi;\Phi \vdash_{\mathcal{E}} \tau \equiv \sigma$;*
- *τ is $(p, \sigma^{\neg p}, \mathcal{A})$-specified and σ is $(\neg p, \tau^p, \mathcal{A})$-specified.*

We also need a generalisation of the dereliction algorithm, where any index I can be used instead of 0 in DER. It is necessary for type checking a fix point.

Lemma 4.6 *There is an algorithm* SUBS *such that, for any type τ primitive for ϕ and σ primitive for (ϕ, a) that have the same skeleton,* $\mathrm{SUBS}((\sigma, \tau); a; \mathrm{I}; \phi; \Phi; p) = (\mathcal{A}, \mathscr{R})$ *where,*
- *for any $\mathcal{E} \supseteq \mathcal{A}$, $\mathcal{E} \Vdash \bigwedge \mathscr{R} \iff \phi;\Phi \vdash_{\mathcal{E}} \sigma\{\mathrm{I}/a\} \equiv \tau$;*
- *σ is $(p, \tau^p, \mathcal{A})$-specified and τ is $(\neg p, \sigma^{\neg p}, \mathcal{A})$-specified.*

The following lemma is needed to type check the (If) rule. It requires to use a function symbol $\mathrm{if}(a,b,c)$ with two equations: $\mathrm{if}(0,b,c) = b$ and $\mathrm{if}(a+1,b,c) = c$.

Lemma 4.7 *There is an algorithm* FORK *such that for any types τ_1, τ_2 with the same skeleton, that are primitive for ϕ,*
$\mathrm{FORK}(\tau_1; \tau_2; \mathrm{I}; \phi; \Phi; p) = (\sigma, \mathcal{A}, \mathscr{R})$ *where,*
- *σ is primitive for (a, ϕ), and $(\![\sigma]\!) = (\![\tau_i]\!)$;*
- *For any equational program $\mathcal{E} \supseteq \mathcal{A}$, $\mathcal{E} \Vdash \bigwedge \mathscr{R}$ iff*
  $\phi;\Phi, \mathrm{I} = 0 \vdash_{\mathcal{E}} \tau_1 \equiv \sigma$ and $\phi;\Phi, \mathrm{I} \geqslant 1 \vdash_{\mathcal{E}} \tau_2 \equiv \sigma$;
- *σ is $(p, \tau_1^p \cup \tau_2^p, \mathcal{A})$-specified, and τ_1 and τ_2 are $(\neg p, \sigma^{\neg p}, \mathcal{A})$-specified.*

4.5 The Type Inference Procedure

In this section, we will describe the core of our type inference algorithm. This consists in a recursive algorithm GEN which decorates a PCF type derivation d_{PCF}, producing in output a dℓPCF judgement, together with an equational program and a set of side conditions. In order to correctly create fresh symbols and to format side conditions properly, the main recursive function GEN also receives a set of index variables ϕ and a set of constraints Φ in input. Thus, it has the following signature:

$$\mathrm{GEN}(\phi, \Phi, d_{\mathsf{PCF}}) = (\Gamma \vdash_{\mathrm{I}} t : \tau; \mathcal{A}; \mathscr{R}).$$

We will prove that the the output of GEN satisfies the following two invariants:
- *Decoration.* $(\![\Gamma]\!) \vdash t : (\![\tau]\!)$ is the judgement concluding d_{PCF}.
- *Polarity.* $\phi;\Phi;\Gamma \vdash_{\mathrm{I}}^{\mathcal{A}} t : \tau$ is correctly specified (Definition 4.5).

The algorithm GEN proceeds by inspecting d_{PCF} in an inductive manner. It first annotates the types in the conclusion judgement with fresh function symbols to get a dℓPCF judgement $\mathcal{J}$. Then a recursive call is performed on the immediate sub-derivations of d_{PCF}, this way obtaining some dℓPCF typing judgement $\mathcal{J}_i$. Finally GEN generates, calling the auxiliary algorithms, the equations on function symbols that allow to derive $\mathcal{J}$ from the $\mathcal{J}_i$'s, The equations are written in $\mathcal{A}$, and the required assumptions of index convergence in $\mathscr{R}$.

Decoration and Polarity are the invariants of the algorithm GEN. In particular, the auxiliary algorithms are always called with the appropriate parameters, this way enforcing Polarity.

The algorithm computing GEN proceeds by case analysis on d_{PCF}.
- Assume $d_{\mathsf{PCF}} = \dfrac{}{y_1 : U_1, \ldots, y_k : U_k, x : T \vdash x : T}(Ax)$.
  For each i, let $B_i = \alpha(\phi; U_i)$ and $\mathcal{A}_i = \mathrm{WEAK}(B_i; \phi; b_i; -)$ (where all the b_i's are fresh). Then let σ and τ produced by $\alpha(\phi; T)$, and write $(\mathcal{A}_0; \mathscr{R}) = \mathrm{EQ}((\tau; \sigma); \phi; \Phi; +)$. Thus return $(\Gamma, x : \sigma \vdash_0 x : \tau; \mathcal{A}; \mathscr{R})$ with

19

$$\begin{cases} \mathcal{A} &= \mathcal{A}_0 \cup \bigcup_{i \leqslant k}(\mathcal{A}_i \cup \{h_i := 0\}) \\ \Gamma &= \{y_i : [b_i < h_i] \cdot B_i\}_{i \leqslant k} \end{cases}$$

where the h_i's are fresh.

- Suppose $d_{\mathsf{PCF}} = \dfrac{}{y_1 : U_1, \ldots, y_k : U_k \vdash \underline{\mathtt{n}} : \mathtt{Nat}}(n)$.

  Again, let $B_i = \alpha(\phi; U_i)$ and $\mathcal{A}_i = \mathrm{WEAK}(B_i; \phi; b_i; -)$. Let $i(\phi)$ be a fresh function symbol. Then return $(\Gamma \vdash_0 \underline{\mathtt{n}} : \mathtt{Nat}[i(\phi)]; \mathcal{A}; \varnothing)$ where the h_i's are fresh symbols and

  $$\mathcal{A} = \bigcup_i (\mathcal{A}_i \cup \{h_i(\phi) := 0\}) \cup \{ i(\phi) := \underline{n} \};$$

  $$\Gamma = \{y_i : [b_i < h_i(\phi)] \cdot B_i\}_{i \leqslant k}.$$

  $$\vdots$$

- If $d_{\mathsf{PCF}} = \dfrac{(d'_{\mathsf{PCF}}): \ \Pi \vdash t : \mathtt{Nat}}{\Pi \vdash \mathtt{s}(t) : \mathtt{Nat}}(s)$.

  Let $(\Gamma \vdash_{\mathrm{K}} t : \mathtt{Nat}[j(\phi)]; \mathcal{A}_0; \mathscr{R}) = \mathrm{GEN}(d'_{\mathsf{PCF}}, \phi, \Phi)$, and let $i(\phi)$ be a fresh symbol. Then return $(\Gamma \vdash_{\mathrm{K}} \mathtt{s}(t) : \mathtt{Nat}[i(\phi)]; \mathcal{A}; \mathscr{R})$ where $\mathcal{A} = \mathcal{A}_0 \cup \{ i(\phi) := j(\phi) + 1 \}$.

- For the typing rule (p), we do the same as previously but with $i(\phi) = j(\phi) - 1$ instead of $i(\phi) = j(\phi) + 1$ in the equational program.

  $$\vdots \qquad\qquad \vdots \qquad\qquad \vdots$$

- If $d_{\mathsf{PCF}} = \dfrac{(d_{\mathsf{PCF}0}): \ \Pi \vdash t : \mathtt{Nat} \quad (d_{\mathsf{PCF}1}): \ \Pi \vdash u_1 : T \quad (d_{\mathsf{PCF}2}): \ \Pi \vdash u_2 : T}{\Pi \vdash \ \mathtt{ifz}\ t\ \mathtt{then}\ u_1\ \mathtt{else}\ u_2 : T}(If)$,

  let $(\Gamma \vdash_{\mathrm{K}} t : \mathtt{Nat}[j(\phi)]; \mathcal{A}_0; \mathscr{R}_0) = \mathrm{GEN}(d_{\mathsf{PCF}0}, \phi, \Phi)$, and $(\Delta_i \vdash_{\mathrm{H}_i} u_i : \tau_i; \mathcal{A}_i; \mathscr{R}_i) = \mathrm{GEN}(d_{\mathsf{PCF}i}, \phi, \Phi_i)$ (for $i = 1, 2$), with $\Phi_1 = (\Phi, j(\phi) = 0)$ and $\Phi_2 = (\Phi, j(\phi) \geqslant 1)$. τ_1 and τ_2 both have skeleton T and are primitive for ϕ (since Decoration and Polarity hold by induction hypothesis). So we can compute $(\tau'; \mathcal{A}'; \mathscr{R}') = \mathrm{FORK}(\tau_1; \tau_2; j(\phi); \phi; \Phi; +)$.
  In the same way, Δ_1 and Δ_2 have the same skeleton Π. So for each $(y_i : \sigma_i)$ in Δ_1 there is $(y_i : \sigma'_i)$ in Δ_2 so that we can compute $(\sigma''_i, \mathcal{A}'_i, \mathscr{R}'_i) = \mathrm{FORK}(\sigma_i; \sigma'_i; j(\phi); \phi; \Phi; -)$. Let $\Delta = \{y_i : \sigma''_i\}$ and write $[a_i < m_i] \cdot C_i = \sigma''_i$. Δ also has skeleton Π, and so has Γ. So for each i, there is some variable declaration $y_i : [a_i < n_i] \cdot B_i$ in Γ, so that we can compute $(A_i, \mathcal{A}''_i, \mathscr{R}''_i) = \mathrm{CTR}((B_i, C_i, n_i); (m_i, a_i); \phi; \Phi; -;)$.
  Then let $h_i(\phi)$ be a fresh symbol for each i, and return $(\Gamma' \vdash_{\mathrm{K}+\mathrm{H}} \ \mathtt{ifz}\ t\ \mathtt{then}\ u_1\ \mathtt{else}\ u_2 : \tau'; \mathcal{A}; \mathscr{R})$ where

  $$\begin{aligned} \mathscr{R} &= \mathscr{R}_0 \cup \mathscr{R}_1 \cup \mathscr{R}_2 \cup \mathscr{R}' \cup \bigcup_i(\mathscr{R}'_i \cup \mathscr{R}''_i) \\ \mathcal{A} &= \mathcal{A}_0 \cup \mathcal{A}_1 \cup \mathcal{A}_2 \cup \mathcal{A}' \cup \{ h_i(\phi) := n_i + m_i \}_i \\ &\quad \cup \bigcup_i (\mathcal{A}'_i \cup \mathcal{A}''_i) \\ \Gamma' &= \{y_i : [a_i < h_i(\phi)] \cdot A_i\}_i \\ \mathrm{H} &= \mathrm{if}(j(\phi), \mathrm{H}_1, \mathrm{H}_2) \end{aligned}$$

- If $d_{\mathsf{PCF}} =$

  $$\dfrac{d^1_{\mathsf{PCF}} : \Pi \vdash t : U \Rightarrow T \quad d^2_{\mathsf{PCF}} : \Pi \vdash u : U}{\Pi \vdash t\ u : T}$$

  let $(\Gamma_1 \vdash_{\mathrm{K}} t : [a < f(\phi)] \cdot (\sigma_1 \multimap \sigma_2); \mathcal{A}_1; \mathscr{R}_1) = \mathrm{GEN}(\phi; \Phi; d^1_{\mathsf{PCF}})$, and $(\Gamma_2 \vdash_{\mathrm{H}} u : \tau; \mathcal{A}_2; \mathscr{R}_2) = \mathrm{GEN}(\phi; \Phi; d^2_{\mathsf{PCF}})$. Let $(\mathcal{B}, \mathscr{S}) = \mathrm{DER}((\sigma_1, \tau); a; \phi; \Phi; +)$. We then annotate T: let $\tau_2 = \alpha(\phi; T)$, and let $(\mathcal{C}, \mathscr{U}) = \mathrm{DER}((\sigma_2, \tau_2); a; \phi; \Phi; -)$. Then we build a context equivalent to $\Gamma_1 \uplus \Gamma_2$: for any $y : [b_y < i_y(\phi)] \cdot B_y$ in Γ_1, there is some $y : [b_y < j_y(\phi)] \cdot C_y$ in Γ_2 (possibly after some α-conversion). Then let $A_y = \alpha((b_y, \phi); (\![B_y]\!))$, and $(\mathcal{A}_y; \mathscr{R}_y) = \mathrm{CTR}((A_y, B_y, C_y); (i_y(\phi), j_y(\phi)); b_y; \phi; \Phi; -)$. There are $\Delta_i \equiv \Gamma_i$ (for $i = 1, 2$) such that $\phi; \Phi \vdash_{\mathcal{E}} \{y : [b_y < i_y(\phi) + j_y(\phi)] \cdot A_y\}_y \equiv \Delta_1 \uplus \Delta_2$, for every $\mathcal{E} \supseteq \bigcup_y \mathcal{A}_y$ such that

$\mathcal{E} \Vdash \bigwedge(\bigcup_y \mathscr{R}_y)$. Thus let h_y's be fresh symbols and return $(\Delta \vdash_{K+H} t\,u : \tau_2; \mathcal{A}; \mathscr{R})$ with

$$\mathscr{R} = \mathscr{R}_1 \cup \mathscr{R}_2 \cup \mathscr{S} \cup \mathscr{U}$$
$$\cup \bigcup_y (\mathscr{R}_y \cup \{\phi; \Phi \models i_y(\phi) + j_y(\phi) \downarrow\});$$
$$\mathcal{A} = \mathcal{A}_1 \cup \mathcal{A}_2 \cup \mathcal{B} \cup \mathcal{C} \cup \{f(\phi) := 1\}$$
$$\cup \bigcup_y (\mathcal{A}_y \cup \{h_y(\phi) := i_y(\phi) + j_y(\phi)\});$$
$$\Delta = \{y : [b_y < h_y(\phi)] \cdot A_y\}_y.$$

- Assume that d_{PCF} is

$$\frac{d'_{\mathsf{PCF}} : \Pi, x : U \vdash t : T}{\Pi \vdash \lambda x.t : U \Rightarrow T}$$

Let a be a fresh index variable, and $i(\phi)$ be a fresh function symbol, and compute $(\Gamma, x : \sigma \vdash_K t : \tau; \mathcal{B}; \mathscr{S}) = \text{GEN}(d'_{\mathsf{PCF}}, (a, \phi), (a < i(\phi), \Phi))$. We build a context equivalent to $\sum_{a < i(\phi)} \Gamma$: for every $y : [b_y < j_y(a, \phi)] \cdot B_y \in \Gamma$, let $A_y = \alpha((b_y, \phi); (\![B_y]\!))$, let $h_y(\phi)$ be a fresh symbol, and write $(\mathcal{A}_y, \mathscr{R}_y) = \text{DIG}((A_y, B_y); (i(\phi), j_y(a, \phi)); \phi; (a, b_y); \Phi; -)$. Then return

$$(\Delta \vdash_{i(\phi) + \sum_{a < i(\phi)} K} \lambda x.t : [a < i(\phi)] \cdot (\sigma \multimap \tau); \mathcal{A}; \mathscr{R})$$

where

$$\mathscr{R} = \mathscr{S} \cup \bigcup_y (\mathscr{R}_y \cup \{\phi; \Phi \models_{\mathcal{E}} \sum_{a < i(\phi)} j_y(a, \phi) \downarrow\});$$
$$\mathcal{A} = \mathcal{B} \cup \cup_y (\mathcal{A}_y \cup \{h_y(\phi) := \sum_{a < i(\phi)} j_y(a, \phi)\});$$
$$\Delta = \{y : [b_y < h_y(\phi)] \cdot A_y\}_y.$$

$$\vdots$$

- If $d_{\mathsf{PCF}} = \dfrac{(d'_{\mathsf{PCF}}) : \ \Pi, x : T \vdash t : T}{\Pi \vdash \mathtt{fix}\ x.t : T}(Fix)$.

Let b be a fresh index variable, and $h(\phi)$ be a fresh function symbol, and compute $(\Gamma, x : [a < i(b, \phi)] \cdot A \vdash_J t : [a < m(b, \phi)] \cdot B; \mathcal{A}'; \mathscr{R}') = \text{GEN}(d'_{\mathsf{PCF}}, (b, \phi), (b < h(\phi), \Phi))$ (up to α-conversion we can assume that the index variable a is the same in the context and the type). Because of the type equations we have to ensure (namely $(a, b, \phi); (a < i(b, \phi), b < h(\phi), \Phi) \vdash_{\mathcal{E}} B\{0/a\}\{1 + b + \bigcirc_b^{b+1,a} i(b, \phi)/b\} \equiv A$), we will use intermediate types.

First we annotate T: let $\sigma = \alpha(b, \phi; T)$. Up to α-conversion, $\sigma = [a < l(b, \phi)] \cdot C$ for some C. Let $(\mathcal{A}_1, \mathscr{R}_1) = \text{DER}((B, C); a; (b, \phi); b < h(\phi), \Phi; -)$. By Lemma 4.1, C is $(+, B^+, \mathcal{A}_1)$-specified and B is $(-, C^-, \mathcal{A}_1)$-specified, and for any $\mathcal{E} \supseteq \mathcal{A}_1$,

$$\mathcal{E} \Vdash \bigwedge \mathscr{R}_1 \quad \Longleftrightarrow \quad b, \phi; b < h(\phi), \Phi \vdash_{\mathcal{E}} B\{0/a\} \equiv C .$$

Now let $(\mathcal{A}_2, \mathscr{R}_2) = \text{SUBS}((C, A); b; 1 + b + \bigcirc_b^{b+1,a} i(b, \phi); (a, b, \phi); a < i(b, \phi), b < h(\phi), \Phi; -)$. By Lemma 4.6, A is $(+, C^+, \mathcal{A}_2)$-specified and C is $(-, A^-, \mathcal{A}_2)$-specified, and for any $\mathcal{E} \supseteq \mathcal{A}_2$, $\mathcal{E} \Vdash \bigwedge \mathscr{R}_2$ iff

$$a, b, \phi; a < i(b, \phi), b < h(\phi), \Phi \vdash_{\mathcal{E}} C\{1 + b + \overset{b+1,a}{\underset{b}{\bigcirc}} i(b, \phi)/b\} \equiv A .$$

Here

Hence $\mathcal{E}_1 \cup \mathcal{E}_2$ is positively defined for A and negatively for B, and $\mathcal{E}_1 \cup \mathcal{E}_2 \subseteq \mathcal{F} \models (\mathscr{R}_1, \mathscr{R}_2)$ implies

$$a, b, \phi; a < \mathrm{i}(b, \phi), b < \mathrm{h}(\phi), \Phi \vdash_{\mathcal{F}} B\{0/a\}\{1 + b + \bigotimes_{b}^{b+1,a} \mathrm{i}(b, \phi)/b\} \equiv A.$$

Again we annotate T: let $\tau = \alpha(\phi; T)$. Up to α-conversion, $\tau = [a < \mathrm{k}(\phi)] \cdot D$ for some D. Let $(\mathcal{E}_3, \mathscr{R}_3) = \mathrm{SUBS}((C, D); b; \bigotimes_b^{0,a} \mathrm{i}(b, \phi); (a, \phi); (a < \mathrm{k}, \Phi); -)$. Then $\mathcal{E}_3$ is positively specified for D and negatively for A, and if $\mathcal{E}_3 \subseteq \mathcal{F} \models \mathscr{R}_3$ then

$$a, \phi; a < \mathrm{k}, \Phi \vdash_{\mathcal{F}} C\{\bigotimes_b^{0,a} \mathrm{i}(b, \phi)/b\} \equiv D.$$

Hence $\mathcal{E}_1 \cup \mathcal{E}_2 \cup \mathcal{E}_3$ is positively defined for D and negatively for B, and for any $\mathcal{F} \supseteq \mathcal{E}_1 \cup \mathcal{E}_2 \cup \mathcal{E}_3$, $\mathcal{F} \models (\mathscr{R}_1, \mathscr{R}_2, \mathscr{R}_3)$ implies $a, \phi; a < \mathrm{k}, \bigotimes_b^{0,a} \mathrm{i}(b, \phi) < \mathrm{h}(\phi), \Phi \vdash_{\mathcal{F}} B\{0/a\}\{\bigotimes_b^{0,a} \mathrm{i}(b, \phi)/b\} \equiv D$. Hence for any $\mathcal{F} \supseteq \mathcal{E}_1 \cup \mathcal{E}_2 \cup \mathcal{E}_3 \cup \{\mathrm{h}(\phi) = \bigotimes_b^{0,\mathrm{k}(\phi)} \mathrm{i}(b, \phi)\}$ such that $\mathcal{F} \models (\mathscr{R}_1, \mathscr{R}_2, \mathscr{R}_3)$ and $\phi; \Phi \models_{\mathcal{F}} \bigotimes_b^{0,\mathrm{k}(\phi)} \mathrm{i}(b, \phi) \downarrow$,

$$\phi; \Phi \vdash_{\mathcal{F}} [a < \mathrm{k}] \cdot B\{0/a\}\{\bigotimes_b^{0,a} \mathrm{i}(b, \phi)/b\} \equiv [a < \mathrm{k}] \cdot D.$$

To deal with the context, we proceed as for $(\multimap)$: for every $y : [b_y < \mathrm{J}_y] \cdot B_y \in \Gamma$, let $(A_y, \mathcal{E}_y, \mathscr{R}_y) = \mathrm{DIG}((B_y, \mathrm{h}(\phi)); (\mathrm{J}_y, \phi); b; (b_y, \Phi); -;)$. By Lemma 4.3, $\mathcal{E}_y$ is positively specified for B_y and negatively for A_y, and $\mathcal{E}_y \subseteq \mathcal{F} \models \mathscr{R}_y$ implies

$$\phi; \Phi \vdash_{\mathcal{F}} [b_y < \sum_{b<\mathrm{h}(\phi)} \mathrm{J}_y] \cdot A_y \equiv \sum_{b<\mathrm{h}(\phi)} [b_y < \mathrm{J}_y] \cdot B_y'$$

for some B_y' such that $\phi; \Phi, b < \mathrm{h}(\phi), b_y < \mathrm{J}_y \vdash_{\mathcal{F}} B_y' \equiv B_y$. Let $\mathrm{l}_y(\phi)$ be a fresh symbol, and write $\mathcal{E}_y' = \mathcal{E}_y \cup \{\mathrm{l}_y(\phi) = \sum_{b<\mathrm{h}(\phi)} \mathrm{J}_y\}$ and $\mathscr{R}_y' = \mathscr{R}_y \cup \{\phi; \Phi \models \sum_{b<\mathrm{l}(\phi)} \mathrm{J}_y \downarrow\}$.
Finally we return $(d_\mathsf{v}, \mathscr{R})$ where

$$\mathscr{R} = \mathscr{R}' \cup \mathscr{R}_1 \cup \mathscr{R}_2 \cup \mathscr{R}_3 \cup \{\phi; \Phi \models \bigotimes_b^{0,\mathrm{k}(\phi)} \mathrm{i}(b, \phi) \downarrow\} \cup \bigcup_y \mathscr{R}_y'$$
$$d_\mathsf{v} = \frac{d_\mathsf{v}'}{\phi; \Phi; \Delta \vdash_{\mathrm{h}(\phi)+\sum_{b<\mathrm{h}(\phi)} \mathrm{J}}^{\mathcal{E}} \mathtt{fix}\ x.t : \tau}$$
$$\mathcal{E} = \mathcal{E}' \cup \mathcal{E}_1 \cup \mathcal{E}_2 \cup \mathcal{E}_3 \cup \{\mathrm{h}(\phi) = \bigotimes_b^{0,\mathrm{k}(\phi)} \mathrm{i}(b, \phi)\} \cup \bigcup_y \mathcal{E}_y'$$
$$\Delta = \{y : [b_y < \mathrm{l}_y(\phi)] \cdot A_y\}_y$$

Lemma 4.8 *For every ϕ, Φ, and every PCF derivation d_{PCF}, $\mathrm{GEN}(\phi; \Phi; d_{\mathsf{PCF}})$ is well defined on the form $(\Gamma \vdash_{\mathrm{I}} t : \tau; \mathcal{A}; \mathscr{R})$, and satisfies Decoration and Polarity.*

4.6 Correctness

The algorithm we have just finished describing needs to be proved sound and complete with respect to dℓPCF$_\mathsf{V}$ typing. As usual, this is not a trivial task. Moreover, linear dependent types have a semantic nature which makes the task of formulating (if not proving) the desired results even more challenging.

4.6.1 Soundness

A type inference procedure is *sound* when the inferred type can actually be derived by way of the type system at hand. As already remarked, GEN outputs an equational program $\mathcal{A}$ which possibly contains unspecified symbols and which, as a consequence, cannot be exploited in typing. Moreover, the role of the set of proof obligations in $\mathscr{R}$ is maybe not clear at first. Actually, soundness holds for *every* completely specified $\mathcal{E} \supseteq \mathcal{A}$ which makes the proof obligations in $\mathscr{R}$ true:

Theorem 4.9 (Soundness) *If d_{PCF} is a PCF derivation for t, then for any ϕ and Φ, $\mathrm{GEN}(\phi; \Phi; d_{PCF}) = (\Gamma \vdash_{\mathrm{I}} t : \tau; \mathcal{A}; \mathscr{R})$ where $\phi; \Phi; \Gamma \vdash_{\mathrm{I}}^{\mathcal{A}} t : \tau$ is correctly specified and for any $\mathcal{E} \supseteq \mathcal{A}$,*

$$\mathcal{E} \Vdash \bigwedge \mathscr{R} \quad \Longrightarrow \quad \phi; \Phi; \Gamma \vdash_{\mathrm{I}}^{\mathcal{E}} t : \tau \text{ is derivable and precise.}$$

4.6.2 Completeness

But are we sure that *at least* one type derivation can be built from the outcome of GEN *if* one such type derivation exists? Again, it is nontrivial to formulate the fact that this is actually the case.

Theorem 4.10 (Completeness) *If $\phi; \Phi; \Delta \vdash_{\mathrm{J}}^{\mathcal{E}} t : \sigma$ is a precise $\mathsf{d}\ell\mathsf{PCF_V}$ judgement derivable with d_{v}, then $\mathrm{GEN}(\phi; \Phi; [\![d_{\mathsf{v}}]\!])$ is of the form $(\Gamma \vdash_{\mathrm{I}} t : \tau; \mathcal{A}; \mathscr{R})$, and there is $\mu : \mathcal{E} \models \mathcal{A}$ such that $\mu \Vdash \bigwedge \mathscr{R}$.*

A direct consequence of soundness and completeness (and the remark on Definition 4.3) is the following:

Corollary 4.11 *If a closed term t is typable in PCF with type Nat and a derivation d_{PCF}, then*

$$\mathrm{GEN}(\varnothing; \varnothing; d_{PCF}) = (\vdash_{\mathrm{I}} t : \mathtt{Nat}[\mathsf{f}]; \mathcal{E}; \mathscr{R})$$

and t is typable in $\mathsf{d}\ell\mathsf{PCF_V}$ iff $\mathcal{E} \Vdash \bigwedge \mathscr{R}$.

5 Type Inference at Work

The type inference algorithm presented in the previous section has been implemented in OCAML[2]. Programs, types, equational programs and side conditions become values of appropriately defined inductive data structures in OCAML, and the functional nature of the latter makes the implementation effort easier. This section is devoted to discussing the main issues we have faced along the process, which is still ongoing.

The core of our implementation is an OCAML function called `CheckBound`. Taking a closed term t having PCF type T in input, `CheckBound` returns a typing derivation d_{v}, an equational program $\mathcal{E}$ and a set of side conditions $\mathscr{R}$. The conclusion of d_{v} is a $\mathsf{d}\ell\mathsf{PCF}$ typing judgement for the input term. If T is a first-order type, then the produced judgement is derivable *iff* all the side conditions in $\mathscr{R}$ are valid (see Corollary 4.11). To do so, `CheckBound` calls (an implementation of) GEN on t and a context ϕ consisting of n unconstrained index variables, where n is the arity of t. This way, `CheckBound` obtains $\mathcal{A}$ and $\mathscr{R}$ as results, and then proceeds as follows:

- If T is Nat, then $\mathcal{A}$ is already completely specified and Corollary 4.11 ensures that we already have what we need.
- If T has a strictly positive arity, then some of the symbols in $\mathcal{A}$ are unspecified, and appropriate equations for them need to be added to $\mathcal{A}$. Take for instance a term s of type $\mathtt{Nat} \Rightarrow \mathtt{Nat}$. `CheckBound`$(s)$ returns $\mathcal{A}$, $\mathscr{R}$, and a typing judgement on the form

$$a; \varnothing \vdash_{\mathrm{K}}^{\mathcal{A}} t : \mathtt{Nat}[\mathsf{g}(a,b)] \overset{b < \mathsf{f}(a)}{\multimap} \mathtt{Nat}[\mathsf{j}(a,b)],$$

  where j is a positive symbol while f and g are negative, thus unspecified in $\mathcal{A}$. $\mathcal{A}$ can be appropriately "completed" by adding the equations $\mathsf{f}(a) := \underline{1}$ and $\mathsf{g}(a,b) := a$ to it. This way, we are insisting on the behaviour of t when fed with *any* natural number (represented by a) and when the environment needs t only once.

How about complexity analysis? Actually, we are already there: the problem of proving the number of machine reduction steps needed by t to be at most $p : \mathbb{N} \to \mathbb{N}$ (where p is, e.g. a polynomial) becomes the problem of checking $\mathcal{E} \Vdash \bigwedge \mathscr{S}$ where $\mathcal{E}$ is the appropriate completion of $\mathcal{A}$, and $\mathscr{S}$ is $\mathscr{R} \cup \{(\mathrm{K}+1)(|t|+2) \leqslant p(a)\}$ (Proposition 3.2).

---

[2]the source code is available at `http://lideal.cs.unibo.it`.

Simplifying Equations.   Equational programs obtained in output from `CheckBound` contains many equations which are trivial (such as $f(a) = \underline{n}$ or $f(a) = g(a)$), and as such can be eliminated. Moreover, instances of forest cardinalities and bounded sums can sometime be greatly simplified. As an example, $\sum_{a<\underline{0}} J$ can always be replaced by $\underline{0}$. This allows, in particular, to turn $\mathcal{A}$ into a set of fewer and simpler rules, thus facilitating the next phase.

A basic simplification procedure has already been implemented, and are called by `CheckBound` on the output of GEN. However, automatically treat the equational program by an appropriate prover would be of course desirable. For this purpose, the possibility for `CheckBound` to interact with MAUDE [8], a system supporting equational and rewriting logic specification, is currently investigated.

Checking Side Conditions.   As already stressed, once `CheckBound` has produced a pair $(\mathcal{A}, \mathcal{R})$, the task we started from, namely complexity analysis of t, is not finished, yet: checking proof obligations in $\mathcal{R}$ is as undecidable as analysing the complexity of t directly, since most of the obligations in $\mathcal{R}$ are termination statements anyway. There is an important difference, however: statements in $\mathcal{R}$ are written in a language (the first-order equational logic) which is more amenable to be treated by already existing automatic and semi-automatic tools.

Actually, the best method would be to first call as many existing automatic provers on the set of side conditions, then asking the programmer to check those which cannot be proved automatically by way of an interactive theorem prover. For this purpose, we have implemented an algorithm translating a pair in the form $(\mathcal{A}, \mathcal{R})$ into a WHY3 [6] theory.   Indeed, WHY3 is an intermediate tool between first order logic and various theorem provers, from SMT solvers to the COQ proof assistant [5][3] on the side conditions produced by `CheckBound`(t). Most of them are actually proved automatically, at least in the few example programs we have mentioned in the course of this paper [4]. Actually, the symbol names used in the WHY3 theory are the same as the ones used for the annotation of the type derivation (that can be printed by `CheckBound`). Hence trying to check interactively some side conditions, the programmer can access both the definition of a symbol in the equational program, and the subterm of t it refers to in the type derivation.

However, using the proof assistant COQ on the few side conditions that remain to be interactively checked *through* the WHY3 tool is not as simple as it should be. This is due to the way we express bounded sum and forest cardinality indexes in a first order logic. To facilitate the work of the programmer checking the side conditions, it would thus be suitable to also translate them directly into COQ, making use of its higher order definition facilities. The interactive theorem prover would be called *directly*, but only on those side conditions that cannot be proved automatically. This of course requires some special care: we would like to preserve a formal link between the WHY3 theory and the statement which that have to be proved in COQ (or in any interactive theorem prover).

All these issues are currently investigated and developed within LIDEAL [25].

6   Related Work

Complexity analysis of higher-order programs has been the object of study of much research. We can for example mention the many proposals for type systems for the λ-calculus which have been shown to correspond in an *extensional sense* to, e.g. polynomial time computable functions as in implicit computational complexity. Many of them can be seen as static analysis methodologies: once a program is assigned a type, an upper bound to its time complexity is relatively easy to be synthesised. The problem with these systems, however, is that they are usually very weak from an *intentional* point of view, since the class of typable programs is quite restricted compared to the class of all terms working within the prescribed resource bounds.

More powerful static analysis methodologies can actually be devised. All of them, however, are limited either to very specific forms of resource bounds or to a peculiar form of higher-order functions or else they

---

[3] ALT-ERGO, CVC3, E-PROVER, GAPPA, SIMPLIFY, SPASS, VAMPIRE, VERIT, YICEX and Z3
[4] The results of these tests are available at `http://lideal.cs.unibo.it/`

do not get rid of higher-order as the underlying logic. Consider, as an example, one of the earliest work in this direction, namely Sands's system of cost closures [28]: the class of programs that can be handled includes the full lazy λ-calculus, but the way complexity is reasoned about remains genuinely higher-order, being based on closures and contexts. In Benzinger's framework [4] higher-order programs are translated into higher-order equations, and the latter are turned into first-order ones; both steps, and in particular the second one, are not completeness-preserving. Recent works on amortised resource analysis are either limited to first-order programs [20] or to linear bounds [23]. A recent proposal by Amadio and Régis-Gianas [2] allows to reason on the the cost of higher-order functional programs by way of so-called cost-annotations, being sure that the actual behaviour of compiled code somehow reflects the annotation. The logic in which cost annotations are written, however, is a form of *higher-order* Hoare logic. None of the proposed systems, on the other hand, are known to be (relatively) complete in the sense we use here.

Ghica's slot games [16] are maybe the work which is closest to ours, among the many in the literature. Slot Games are simply ordinary games in the sense of game semantics, which are however instrumented so as to reflect not only the observable behaviour of (higher-order) programs, but also their performance. Indeed, slot games are fully abstract with respect to an operational theory of improvements due do Sands [29]: this can be seen as the counterpart of our relative completeness theorem. An aspect which has not been investigated much since Ghica's proposal is whether slot games provides a way to perform actual verification of programs, maybe via some form of model checking. As we have already mentioned, linear dependency can be seen as a way to turn games and strategies into types, so one can see the present work also as an attempt to keep programs and strategies closer to each other, this way facilitating verification.

7   Conclusions

A type inference procedure for dℓPCF has been introduced which, given a PCF term, reduces the problem of finding a type derivation for it to the one of solving proof obligations on an equational program, itself part of the output. Truth of the proof obligations correspond to termination of the underlying program. Any type derivation in dℓPCF comes equipped with an expression bounding the complexity of evaluating the underlying program. Noticeably, proof obligations and the related equational program can be obtained in polynomial time in the size of the input PCF program.

The main contribution of this paper consists in having shown that linear dependency is not only a very powerful tool for the precise analysis of higher-order functional programs, but is also a way to effectively and efficiently turn a complex problem (that of evaluating the time complexity of an higher-order program) into a much easier one (that of checking a set of proof obligations for truth).

Although experimental evaluation shows that proof obligations can potentially be handled by modern tools mixing automatic and semi-automatic reasoning, as explained in Section 5, much remains to be done about the technical aspects of turning proof obligations into a form which is suitable to automatic or semi-automatic solving. Actually, many different tools could conceivably be of help here, each of them requiring a specific input format. This implies, however, that the work described here, although not providing a fully-fledged out-of-the-box methodology, has the merit of allowing to factor a complex non-well-understood problem into a much-better-studied problem, namely verification of first-order inequalities on the natural numbers.

References

[1] S. Abramsky, R. Jagadeesan, and P. Malacaria. Full abstraction for PCF. *I & C*, 163(2):409–470, 2000.

[2] R. M. Amadio and Y. Régis-Gianas. Certifying and reasoning on cost annotations of functional programs. *CoRR*, abs/1110.2350, 2011.

[3] G. Barthe, B. Grégoire, and C. Riba. Type-based termination with sized products. In *CSL 2008*, volume 5213 of *LNCS*, pages 493–507. Springer, 2008.

[4] R. Benzinger. Automated higher-order complexity analysis. *Theor. Comput. Sci.*, 318(1-2):79–103, 2004.

[5] Y. Bertot and P. Castéran. *Coq'Art: The Calculus of Inductive Constructions*. Springer, 2004.

[6] F. Bobot, J.-C. Filliatre, C. Marché, and A. Paskevich. Why3: Shepherd your herd of provers. In *First International Workshop on Intermediate Verification Languages*, pages 53–64, 2011.

[7] E. M. Clarke. Programming language constructs for which it is impossible to obtain good hoare axiom systems. *J. ACM*, 26(1):129–147, 1979.

[8] M. Clavel, F. Durán, S. Eker, P. Lincoln, N. Mart´i-Oliet, J. Meseguer, and C. Talcott. The Maude 2.0 system. In *RTA 2003*, volume 2706 of *LNCS*, pages 76–87, 2003.

[9] S. A. Cook. Soundness and completeness of an axiom system for program verification. *SIAM J. on Computing*, 7:70–90, 1978.

[10] K. Crary and S. Weirich. Resource bound certification. In *ACM POPL*, pages 184–198, 2000.

[11] U. Dal Lago. Context semantics, linear logic and computational complexity. In *LICS*, pages 169–178, 2006.

[12] U. Dal Lago and M. Gaboardi. Linear dependent types and relative completeness. In *LICS*, pages 133–142, 2011.

[13] U. Dal Lago and B. Petit. Linear dependent types in a call-by-value scenario. In *ACM PPDP 2012*, pages 115–126, 2012.

[14] E. Denney. Refinement types for specification. In *IFIP-PROCOMET*, pages 148–166, 1998.

[15] M. Felleisen and D. P. Friedman. Control operators, the SECD-machine and the λ-calculus. Technical Report 197, Computer Science Department, Indiana University, 1986.

[16] D. R. Ghica. Slot games: a quantitative model of computation. In *ACM POPL*, pages 85–97, 2005.

[17] J.-Y. Girard. Linear logic. *Theor. Comp. Sci.*, 50:1–102, 1987.

[18] J.-Y. Girard, A. Scedrov, and P. Scott. Bounded linear logic. *Theor. Comp. Sci.*, 97(1):1–66, 1992.

[19] S. Gulwani. Speed: Symbolic complexity bound analysis. In *CAV*, pages 51–62, 2009.

[20] J. Hoffmann, K. Aehlig, and M. Hofmann. Multivariate Amortized Resource Analysis. In *ACM POPL*, pages 357–370, 2011.

[21] J. Hughes, L. Pareto, and A. Sabry. Proving the correctness of reactive systems using sized types. In *ACM POPL*, pages 410–423, 1996.

[22] J. Iversen. An action compiler targeting Standard ML. *Sci. Comput. Program.*, 68(2):79–94, 2007.

[23] S. Jost, K. Hammond, H.-W. Loid, and M. Hofmann. Static determination of quantitative resource usage for higher-order programs. In *ACM POPL*, Madrid, Spain, 2010.

[24] J.-L. Krivine. A call-by-name lambda-calculus machine. *Higher-Order and Symbolic Computation*, 20(3):199–207, 2007.

[25] Lideal. Linear dependant analyser. `http://lideal.cs.unibo.it`, 2012.

[26] J. Maraist, M. Odersky, D. N. Turner, and P. Wadler. Call-by-name, call-by-value, call-by-need and the linear lambda calculus. *Electr. Notes Theor. Comput. Sci.*, 1:370–392, 1995.

[27] G. D. Plotkin. LCF considerd as a programming language. *Theor. Comp. Sci.*, 5:225–255, 1977.

[28] D. Sands. Complexity analysis for a lazy higher-order language. In *ESOP 1990*, volume 432 of *LNCS*, pages 361–376, 1990.

[29] D. Sands. Operational theories of improvement in functional languages (extended abstract). In *Functional Programming*, pages 298–311, 1991.

[30] R. Wilhelm, J. Engblom, A. Ermedahl, N. Holsti, S. Thesing, D. Whalley, G. Bernat, C. Ferdinand, R. Heckmann, T. Mitra, F. Mueller, I. Puaut, P. Puschner, J. Staschulat, and P. Stenstrom. The worst case execution time problem - overview of methods and survey of tools. *ACM Trans. Embed. Comput. Syst.*, 2008.