

AperTO - Archivio Istituzionale Open Access dell'Università di Torino

Type Reconstruction for the Linear π -Calculus with Composite and Equi-Recursive Types

This is a pre print version of the following article:

Original Citation:

Availability:

This version is available <http://hdl.handle.net/2318/153600> since 2017-11-16T12:13:58Z

Publisher:

Springer

Published version:

DOI:10.1007/978-3-642-54830-7_6

Terms of use:

Open Access

Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.

(Article begins on next page)

Type Reconstruction for the Linear π -Calculus with Composite and Equi-Recursive Types

Luca Padovani – Dipartimento di Informatica, Università di Torino, Italy

Abstract. We extend the linear π -calculus with composite and equi-recursive types in a way that enables the sharing of data containing linear values, provided that there is no overlapping access on such values. We show that the extended type system admits a complete type reconstruction algorithm and, as a by-product, we solve the problem of reconstruction for equi-recursive session types.

1 Introduction

The linear π -calculus [11] is a formal model of communicating processes in which channels are either *unlimited* or *linear*. Unlimited channels can be used without restrictions, while linear channels can only be used once for an input/output. Linear channels occur frequently in actual systems, they allow optimisations and efficient implementations [6,5,11], and communications on linear channels enjoy important properties such as interference freedom and partial confluence [13,11].

Type reconstruction is the problem of *inferring* the type of entities – channels in our case – given a program using them. For the linear π -calculus this problem was addressed in [7], although that work did not consider composite or recursive types. The goal of this work is the definition of a type reconstruction algorithm for the linear π -calculus extended with pairs, disjoint sums, and equi-recursive types. These constructs, albeit standard, gain relevance and combine in non-trivial ways with the features of the linear π -calculus. We explain why this is the case in the rest of this section.

By definition, linear channels can only be used for one-shot communications. It is a known fact, however, that more sophisticated interactions can be implemented taking advantage of *channel mobility* using a continuation-passing style [9,2]. The basic idea is that, along with the proper payload of a communication, one can send another channel on which the rest of the conversation takes place. This technique is illustrated below

$$P(x, y) \stackrel{\text{def}}{=} (\nu a)(\bar{x}\langle y, a \rangle \mid P\langle a, y + 1 \rangle) \quad C(x) \stackrel{\text{def}}{=} x(y, z).C\langle z \rangle \quad (1.1)$$

where a *producer* process P sends messages to a *consumer* process C . At each iteration, the producer creates a new channel a , sends the payload y to the consumer on x along with the continuation a on which subsequent communications will take place, and iterates. In parallel, the consumer waits for the payload and the continuation from the producer on x and then iterates. Explicit continuations are key to preserve the order of produced messages. Had we modelled (1.1) re-using the same channel x at each iteration, there would be no guarantee that messages were received in order.

Let us now assign types to the channels in (1.1) starting from x in the consumer. There, x is used once for receiving a pair made of an integer y and another channel z .

Say the type of z is t and note that z is used in $C\langle z \rangle$ in the same place as x , meaning that x and z must have the same type. Then, also x has type t and t must be a channel type satisfying the equation $t = [\mathbf{int} \times t]^{1,0}$ where $\mathbf{int} \times t$ is the type of messages received from the channel and the numbers 1 and 0, henceforth called *uses*, indicate that the channel of type t is used once for input and never for output. Channel x is used once in the producer for sending an integer y and a continuation a . Clearly, the continuation must have type t for that is how it is used in C . Therefore, x in P has type $s = [\mathbf{int} \times t]^{0,1}$ where the uses 0 and 1 indicate that x is never used for input and is used once for output. Finally, there are two (non-binding) occurrences of a in the producer: a in $\bar{x}\langle n, a \rangle$ has type t because that is how a will be used by the consumer; a in $P\langle a, y + 1 \rangle$ has type s because that is how a will be used by the producer at the next iteration. Overall, the uses in the types of a say that a is a linear channel: it is used once by the producer for sending the payload and once by the consumer for receiving it.

Note that linear channels, like a in (1.1), may syntactically occur multiple times and that different occurrences of the same channel may have different yet *compatible* types. In the case of (1.1), the types t and s of the non-binding occurrences of a are compatible because corresponding uses in t and s are never 1 at the same time. The binding occurrence of a in (νa) has type $[\mathbf{int} \times t]^{1,1}$, which can be seen as the *combination* of t and s as defined in [11,14].

One legitimate question is whether and how the notions of type compatibility and combination extend beyond channel types. In this respect, the existing literature lacks definitive and satisfactory answers: the works on (type reconstruction for) the linear π -calculus [11,7] do not consider composite or recursive types. Linear type systems with composite types have been discussed in [5,6] for the linear π -calculus and in [15] for a functional language. These works, however, see linearity as a “contagious” property: every structure that contains linear values becomes linear itself (there are a few exceptions for specific types [10] or relaxed notions of linearity [8]). Such interpretation may be appropriate in a sequential setting, but is not the only sensible one in a concurrent/parallel setting. In fact, it is acceptable (and desirable, for the sake of parallelism) that several processes share the *same* composite data structure, provided that they access *different* linear values stored therein. The problem is whether the type system is accurate enough to capture the fact that there are no overlapping accesses to the same linear values. To illustrate, consider the type t_{list} satisfying the equation

$$t_{list} = \mathbf{unit} \oplus ([\mathbf{int}]^{0,1} \times t_{list})$$

which is the disjoint sum between \mathbf{unit} and the product $[\mathbf{int}]^{0,1} \times t_{list}$ and which can be used for typing lists of linear channels with type $[\mathbf{int}]^{0,1}$. If we follow [15,5,6], an identifier l having type t_{list} can syntactically occur only once in a program, and a process like for instance $Odd\langle l \rangle \mid Even\langle l \rangle$ where l occurs twice is illegal. However, suppose that Odd and $Even$ are the two processes defined by

$$\begin{aligned} Odd(x) &\stackrel{\text{def}}{=} \mathbf{case } x \text{ of } [] \Rightarrow \mathbf{0} & Even(x) &\stackrel{\text{def}}{=} \mathbf{case } x \text{ of } [] \Rightarrow \mathbf{0} \\ & y : z \Rightarrow \bar{y}\langle 3 \rangle \mid Even\langle z \rangle & & y : z \Rightarrow Odd\langle z \rangle \end{aligned}$$

which walk through a list x : if x is the empty list $[]$ they do nothing; if x has head y and tail z , Odd sends 3 on y and continues as $Even\langle z \rangle$ while $Even$ ignores y and continues

as $Odd\langle z \rangle$. So, $Odd\langle l \rangle$ sends 3 on every odd-indexed channel in l and $Even\langle l \rangle$ sends 3 on every even-indexed channel in l . The fact that Odd and $Even$ access different linear values of a list is reflected in the two (mutually recursive) types of x in Odd and $Even$:

$$t_{odd} = \mathbf{unit} \oplus ([\mathbf{int}]^{0,1} \times t_{even}) \quad \text{and} \quad t_{even} = \mathbf{unit} \oplus ([\mathbf{int}]^{0,0} \times t_{odd}) \quad (1.2)$$

where the two 0's in t_{even} denote that $Even$ does not use at all the first (and more generally every odd-indexed) element of its parameter x . The key observation is that just like a was allowed to occur twice in (1.1) with two compatible types t and s whose combination was $[\mathbf{int} \times t]^{1,1}$, then we can allow l to occur twice in $Odd\langle l \rangle \mid Even\langle l \rangle$ given that the two occurrences of l are used according to two compatible types t_{odd} and t_{even} whose combination is t_{list} . The “only” difference is that, while t and s were channel types and their combination could be expressed simply by combining the uses in them, t_{odd} and t_{even} are recursive, structured types that combine to t_{list} in the limit.

To summarise, composite and recursive types are basic yet fundamental features whose smooth integration in the linear π -calculus requires some care. In this work we extend the linear π -calculus with composite and recursive types in such a way that multiple processes can safely share structured data containing linear values and we define a complete type reconstruction algorithm for the extended type system.

We proceed with the formal definition of the language and of the type system (Section 2). The type reconstruction algorithm consists of a constraint generation phase (Section 3) and a constraint solving phase (Section 4). Section 5 concludes. The full version of the paper (with proofs) and a Haskell implementation of the algorithm are available on the author's home page.

2 The Linear π -Calculus

We let m, n, \dots range over integer numbers; we use a countable set of **channels** a, b, \dots and a disjoint countable set of **variables** x, y, \dots ; **names** u, v, \dots are either channels or variables. We work with the asynchronous π -calculus extended with constants, pairs, and disjoint sums. The syntax of expressions and processes is defined below:

$$\begin{aligned} e &::= n \mid u \mid (e, e) \mid \mathbf{inl} \ e \mid \mathbf{inr} \ e \mid \dots \\ P &::= \mathbf{0} \mid u(x).P \mid \bar{u}\langle e \rangle \mid (P \mid Q) \mid *P \mid (va)P \mid \mathbf{let} \ x, y = e \ \mathbf{in} \ P \mid \mathbf{case} \ e \ \mathbf{of} \\ &\quad \{ \mathbf{inl} \ x \Rightarrow P, \\ &\quad \mathbf{inr} \ y \Rightarrow Q \} \end{aligned}$$

Expressions e, \dots are either integers, names, pairs (e_1, e_2) of expressions, or the injection ($i \ e$) of an expression e using the constructor $i \in \{\mathbf{inl}, \mathbf{inr}\}$. **Values** v, w, \dots are expressions without variables.

Processes P, Q, \dots comprise the standard constructs of the asynchronous π -calculus. In addition to these, we include two process forms for deconstructing pairs and disjoint sums. In particular, the process $\mathbf{let} \ x_1, x_2 = e \ \mathbf{in} \ P$ evaluates e , which must result into a pair v_1, v_2 , binds each v_i to x_i , and continues as P .¹ The process $\mathbf{case} \ e \ \mathbf{of} \{i \ x_i \Rightarrow$

¹ As pointed out by one reviewer, this construct is superfluous, because the type system we are about to define allows linear pairs to be accessed more than once with the conventional projection operators. We have added support for pair projections in the implementation, but we keep the \mathbf{let} construct in the formal presentation.

Table 1. Reduction of processes.

[R-COMM] $\bar{a}(v) a(x).Q \xrightarrow{a} Q\{v/x\}$	[R-LET] $\mathbf{let} \ x, y = (v, w) \ \mathbf{in} \ P \xrightarrow{\tau} P\{v, w/x, y\}$	
[R-CASE] $\frac{k \in \{\mathbf{inl}, \mathbf{inr}\}}{\mathbf{case} \ (k \ v) \ \mathbf{of} \ \{i \ x_i \Rightarrow P_i\}_{i=\mathbf{inl}, \mathbf{inr}} \xrightarrow{\tau} P_k\{v/x_k\}}$	[R-PAR] $\frac{P \xrightarrow{\ell} P'}{P Q \xrightarrow{\ell} P' Q}$	
[R-NEW 1] $\frac{P \xrightarrow{a} Q}{(va)P \xrightarrow{\tau} (va)Q}$	[R-NEW 2] $\frac{P \xrightarrow{\ell} Q \quad \ell \neq a}{(va)P \xrightarrow{\ell} (va)Q}$	[R-STRUCT] $\frac{P \equiv P' \quad P' \xrightarrow{\ell} Q' \quad Q' \equiv Q}{P \xrightarrow{\ell} Q}$

$P_i\}_{i=\mathbf{inl}, \mathbf{inr}}$ evaluates e , which must result into a value $(i \ v)$ for $i \in \{\mathbf{inl}, \mathbf{inr}\}$, binds v to x_i and continues as P_i . Notions of **free names** $\text{fn}(P)$ and **bound names** $\text{bn}(P)$ of P are as expected. We identify processes modulo renaming of bound names and we write $P\{v/x\}$ for the capture-avoiding substitution of v for the free occurrences of x in P .

The operational semantics of the language is defined in terms of a structural congruence relation and a reduction relation, as usual. **Structural congruence** \equiv is completely standard (in particular, it includes the law $*P \equiv *P | P$). **Reduction** is defined in Table 1 and is also conventional, except that, like in [11], we decorate the relation with **labels** ℓ that are either channels or the special symbol τ , denoting an unobservable action: in [R-COMM] the label is the channel a on which a message is exchanged, while in [R-LET] and [R-CASE] it is τ to denote the fact that these are internal computations not involving communications. Rules [R-PAR], [R-NEW 1], and [R-NEW 2] propagate labels through parallel compositions and restrictions. In [R-NEW 1], the label a becomes τ when it escapes the scope of a . Rule [R-STRUCT] closes reduction under structural congruence.

The type system makes use of a countable set of **type variables** α, β, \dots and of **uses** κ, \dots which are elements of the set $\{0, 1, \omega\}$. Types t, s, \dots are defined by

$$t ::= \mathbf{int} \mid \alpha \mid t \times t \mid t \oplus t \mid [t]^{\kappa_1, \kappa_2} \mid \mu \alpha. t$$

and include the type of integers \mathbf{int} , products $t_1 \times t_2$ typing values of the form (v_1, v_2) where v_i has type t_i for $i = 1, 2$, and disjoint sums $t_1 \oplus t_2$ typing values of the form $(\mathbf{inl} \ v)$ where v has type t_1 or of the form $(\mathbf{inr} \ v)$ where v has type t_2 . Throughout the paper we let \odot stand for either \times or \oplus . The type $[t]^{\kappa_1, \kappa_2}$ denotes channels for exchanging messages of type t . The uses κ_1 and κ_2 respectively denote how many input and output operations are allowed on the channel: 0 for none, 1 for a single use, and ω for any number of uses. For example: a channel with type $[t]^{1,1}$ must be used once for receiving a message of type t and once for sending a message of type t ; a channel with type $[t]^{0,0}$ cannot be used; a channel with type $[t]^{\omega, \omega}$ can be used any number of times for sending and/or receiving. Type variables and μ operators are used for building recursive types, as usual. Notions of free and bound type variables are as expected. We assume that every bound type variable is guarded by a constructor to avoid meaningless terms such as $\mu \alpha. \alpha$. We identify types modulo renaming of bound type variables and if their infinite

unfoldings are the same (regular) tree [1]. In particular, we let $\mu\alpha.t = t\{\mu\alpha.t/\alpha\}$ where $t\{s/\alpha\}$ is the capture-avoiding substitution of the free occurrences of α in t with s .

We now define some key relations on uses and types. In particular, \leq is the least **partial order** such that $0 \leq \kappa$ and **compatibility** \asymp is the least relation such that

$$0 \asymp \kappa \quad \kappa \asymp 0 \quad \omega \asymp \omega \quad (2.1)$$

In what follow we will write $\kappa_1 < \kappa_2$ if $\kappa_1 \leq \kappa_2$ and $\kappa_1 \neq \kappa_2$ and $\kappa_1 \vee \kappa_2$ for the **least upper bound** of κ_1 and κ_2 , when it is defined. Note that $1 \not\leq \omega$ does *not* hold. Compatibility determines whether the least upper bound of two uses expresses their combination without any loss of precision. For example, $0 \asymp 1$ because the combination of 0 uses and 1 use is $0 \vee 1 = 1$ use. On the contrary, $1 \not\asymp 1$ because there is no 2 use that expresses the combination of 1 and 1 and ω is less precise than 2. Similarly, $1 \not\asymp \omega$ because there is no use expressing the fact that a channel is used *at least* once.

Every binary relation \mathcal{R}_{use} on uses induces a corresponding relation \mathcal{R}_{type} on types, defined coinductively by the following rules:

$$\text{int } \mathcal{R}_{type} \text{ int} \quad \frac{\kappa_1 \mathcal{R}_{use} \kappa_3 \quad \kappa_2 \mathcal{R}_{use} \kappa_4}{[t]^{\kappa_1, \kappa_2} \mathcal{R}_{type} [t]^{\kappa_3, \kappa_4}} \quad \frac{t_1 \mathcal{R}_{type} s_1 \quad t_2 \mathcal{R}_{type} s_2}{t_1 \odot t_2 \mathcal{R}_{type} s_1 \odot s_2} \quad (2.2)$$

Similarly, the partial operation \vee on uses coinductively induces one on types so that $t \vee s$ is the least upper bound of t and s , if it is defined. Note that \leq is antisymmetric, in particular $t = s$ if and only if $t \leq s$ and $t \geq s$. Note also that $[t]^{\kappa_1, \kappa_2} \mathcal{R}_{type} [s]^{\kappa_3, \kappa_4}$ implies $t = s$ regardless of \mathcal{R}_{type} . The relation $t \asymp t$ is particularly interesting, because it characterises unlimited types, those typing values that must not or need not be used. Linear types, on the other hand, denote values that must be used:

Definition 2.1. We say that t is **unlimited** if $t \asymp t$. We say that t is **linear** otherwise.

Channel types are either limited or unlimited depending on their uses. So, $[t]^{1,0}$, $[t]^{0,1}$, and $[t]^{1,1}$ are all linear types whereas $[t]^{0,0}$ and $[t]^{\omega, \omega}$ are both unlimited. Other types are linear or unlimited according to the channel types occurring in them. For example, $[t]^{0,0} \times [t]^{1,0}$ is linear while $[t]^{0,0} \times [t]^{\omega,0}$ is unlimited. Recursion does not affect the classification of types into linear and unlimited. For example, $\mu\alpha.[\text{int} \times \alpha]^{1,0}$ is a linear type that denotes a channel that must be used once for receiving a pair made of an integer and another channel with the same type.

We type check expressions and processes in **type environments** Γ, \dots , which are finite maps from names to types that we write as $u_1 : t_1, \dots, u_n : t_n$. We identify type environments modulo the order of their bindings, we denote the **empty environment** with \emptyset , we write $\text{dom}(\Gamma)$ for the **domain** of Γ , namely the set of names for which there is a binding in Γ , and Γ_1, Γ_2 for the **union** of Γ_1 and Γ_2 when $\text{dom}(\Gamma_1) \cap \text{dom}(\Gamma_2) = \emptyset$. We extend the relation \asymp between types pointwise to type environments:

$$\begin{aligned} \Gamma_1 + \Gamma_2 &\stackrel{\text{def}}{=} \Gamma_1, \Gamma_2 && \text{if } \text{dom}(\Gamma_1) \cap \text{dom}(\Gamma_2) = \emptyset \\ (\Gamma_1, u : t) + (\Gamma_2, u : s) &\stackrel{\text{def}}{=} (\Gamma_1 + \Gamma_2), u : t \vee s && \text{if } t \asymp s \end{aligned} \quad (2.3)$$

The operator $+$ generalises type combination in [11] and the \uplus operator in [14]. Note that $\Gamma_1 + \Gamma_2$ is undefined if there is $u \in \text{dom}(\Gamma_1) \cap \text{dom}(\Gamma_2)$ and $\Gamma_1(u) \not\asymp \Gamma_2(u)$ and that

Table 2. Type rules for expressions and processes.

Expressions				
$\frac{[\text{T-CONST}]}{\Gamma \asymp \Gamma} \quad \frac{[\text{T-NAME}]}{\Gamma \asymp \Gamma} \quad \frac{[\text{T-PAIR}]}{\Gamma \vdash e : t \quad \Gamma' \vdash e' : s} \quad \frac{[\text{T-INL}]}{\Gamma \vdash e : t} \quad \frac{[\text{T-INR}]}{\Gamma \vdash e : s}$ $\frac{}{\Gamma \vdash n : \mathbf{int}} \quad \frac{}{\Gamma, u : t \vdash u : t} \quad \frac{}{\Gamma + \Gamma' \vdash (e, e') : t \times s} \quad \frac{}{\Gamma \vdash \mathbf{inl} \ e : t \oplus s} \quad \frac{}{\Gamma \vdash \mathbf{inr} \ e : t \oplus s}$				
Processes				
$\frac{[\text{T-IDLE}]}{\Gamma \asymp \Gamma} \quad \frac{[\text{T-IN}]}{\Gamma, x : t \vdash P} \quad \frac{[\text{T-OUT}]}{\Gamma \vdash e : t} \quad \frac{[\text{T-PAR}]}{\Gamma_i \vdash P_i^{(i=1,2)}} \quad \frac{[\text{T-REP}]}{\Gamma \vdash P} \quad \Gamma \asymp \Gamma$ $\frac{}{\Gamma \vdash \mathbf{0}} \quad \frac{}{\Gamma + u : [t]^{\kappa, 0} \vdash u(x).P} \quad \frac{}{\Gamma + u : [t]^{0, \kappa} \vdash \bar{u}(e)} \quad \frac{}{\Gamma_1 + \Gamma_2 \vdash P_1 \mid P_2} \quad \frac{}{\Gamma \vdash *P}$				
$\frac{[\text{T-NEW}]}{\Gamma, a : [t]^{\kappa, \kappa} \vdash P} \quad \frac{[\text{T-LET}]}{\Gamma \vdash e : t \times s} \quad \frac{[\text{T-CASE}]}{\Gamma \vdash e : t \oplus s} \quad \Gamma', x_i : t \vdash P_i^{(i=\mathbf{inl}, \mathbf{inr})}$ $\frac{}{\Gamma \vdash (va)P} \quad \frac{}{\Gamma + \Gamma' \vdash \mathbf{let} \ x, y = e \ \mathbf{in} \ P} \quad \frac{}{\Gamma + \Gamma' \vdash \mathbf{case} \ e \ \mathbf{of} \ \{i \ x_i \Rightarrow P_i\}_{i=\mathbf{inl}, \mathbf{inr}}}$				

$\text{dom}(\Gamma_1 + \Gamma_2) = \text{dom}(\Gamma_1) \cup \text{dom}(\Gamma_2)$. Thinking of type environments as of specifications of the resources used by expressions/processes, $\Gamma_1 + \Gamma_2$ expresses the combined use of the resources specified in Γ_1 and Γ_2 . Any resource occurring in only one of these environments occurs in $\Gamma_1 + \Gamma_2$; any resource occurring in both Γ_1 and Γ_2 must be used according to compatible types, and its type in $\Gamma_1 + \Gamma_2$ is their least upper bound. For example, if a channel is used by a process for sending a message of type \mathbf{int} , it has type $[\mathbf{int}]^{0,1}$ in that process; if it is used by another process for receiving a message of type \mathbf{int} , it has type $[\mathbf{int}]^{1,0}$ in that process. Overall, the two processes in parallel use the channel according to the type $[\mathbf{int}]^{0,1} \vee [\mathbf{int}]^{1,0} = [\mathbf{int}]^{1,1}$.

Type rules for expressions and processes are presented in Table 2. These rules are basically the same as those found in the literature [11,7], and the additional flexibility enabled by our typing discipline is actually a consequence of our notion of type combination \vee . The rules for expressions are unremarkable. Just observe that the part of the type environment that is not used in the expression must be unlimited ($\Gamma \asymp \Gamma$). The idle process does nothing, so it is well typed only in an unlimited environment. Input and output processes require a strictly positive use of the corresponding operation in the type of the channel u used for communication. Rule [T-REP] states that a replicated process $*P$ is well typed in the environment Γ provided that P is well typed in Γ and that Γ is unlimited. The rationale for this is that $*P$ stands for an unbounded number of copies of P composed in parallel, hence it cannot contain (free) linear channels. The rules [T-PAR], [T-LET], and [T-CASE] are conventional. Finally, rule [T-NEW] states that $(va)P$ is well typed if so is P , where P has visibility of a . We require the type of a to have the same uses for input and output. This is not necessary for the soundness of the type system, although it is a sensible choice in practice.

The type system is sound and the results in Section 4.3 of [11] can be formulated in our setting. In particular, the operations on a channel never exceed the uses in its type. It

The novelty is that we need constraints expressing not only the *equality* between types and uses, but also the order \leq and compatibility \succsim , which affect the solution phase in non-trivial ways.

To get started, we generalise uses to *use expressions*, which are either uses or *use variables* ρ, \dots that denote an unknown use. We also define *type expressions* as types without μ 's and where we admit use expressions wherever uses can occur. We keep κ and t for ranging over use and type expressions, respectively, and we say that t is *proper* if it is not a type variable. *Constraints* φ, \dots have one of these forms:

$$\varphi ::= \kappa_1 \mathcal{R}_c \kappa_2 \mid t_1 \mathcal{R}_c t_2$$

where $\mathcal{R} \in \{\leq, <, \succsim, \sim\}$ and \sim is the trivial relation such that $\kappa_1 \sim \kappa_2$ holds for all κ_1 and κ_2 . We call $\kappa_1 \mathcal{R}_c \kappa_2$ a *use constraint* and $t_1 \mathcal{R}_c t_2$ a *type constraint*. The subscript \cdot_c reminds us that $\kappa_1 \mathcal{R}_c \kappa_2$ and $t_1 \mathcal{R}_c t_2$ are just *triples* made of two use or type expressions and a *symbol* \mathcal{R}_c denoting a relation. Equality constraints $=_c$ can be expressed as the conjunction of two constraints \leq_c and \geq_c , given that \leq is antisymmetric. Constraints with the strict order $<_c$ will be generated only for use expressions. Compatibility constraints \succsim_c will be generated for ensuring type and use combination, as well as for expressing the assumption that a type/use is unlimited (see *e.g.* the premise of [T-IDLE]). Finally, \sim_c constraints relate types that must be *structurally coherent*. For example, $[t]^{\kappa_1, \kappa_2} \sim [t]^{\kappa_3, \kappa_4}$ holds regardless of $\kappa_1, \kappa_2, \kappa_3$, and κ_4 (but note that according to (2.2) there must be the same t in the two types). We will see in Section 4 the role of these constraints.

We let \mathcal{C}, \dots range over finite sets of constraints. The *domain* of \mathcal{C} , written $\text{dom}(\mathcal{C})$, is the (finite) set of use and type expressions occurring in the constraints in \mathcal{C} . We let σ range over finite maps from type variables to types and from use variables to uses. The *application* of σ replaces use variables ρ and type variables α with the corresponding uses $\sigma(\rho)$ and types $\sigma(\alpha)$. We write $\sigma\kappa$ and σt for the application of σ to κ and t , respectively. We say that σ is a *solution* of \mathcal{C} if $\sigma t \mathcal{R} \sigma s$ for every $t \mathcal{R}_c s \in \mathcal{C}$ and $\sigma\kappa_1 \mathcal{R} \sigma\kappa_2$ for every $\kappa_1 \mathcal{R}_c \kappa_2 \in \mathcal{C}$. We extend the \leq relation pointwise to solutions and we say that a solution σ for \mathcal{C} is *minimal* if every solution $\sigma' \leq \sigma$ for \mathcal{C} is such that $\sigma \leq \sigma'$. We say that \mathcal{C} is *satisfiable* if it has a solution. We say that \mathcal{C}_1 and \mathcal{C}_2 are *equivalent* if they have the same solutions.

We need two operators for combining and merging type environments in the reconstruction algorithm. They take two type environments Γ_1 and Γ_2 and produce a pair consisting of another type environment Γ and a set of constraints \mathcal{C} :

$$\frac{\text{dom}(\Gamma_1) \cap \text{dom}(\Gamma_2) = \emptyset}{\Gamma_1 \sqcup \Gamma_2 \rightsquigarrow \Gamma_1, \Gamma_2; \emptyset} \quad \frac{\Gamma_1 \sqcup \Gamma_2 \rightsquigarrow \Gamma; \mathcal{C} \quad \alpha \text{ fresh}}{(\Gamma_1, u : t) \sqcup (\Gamma_2, u : s) \rightsquigarrow \Gamma, u : \alpha; \mathcal{C} \cup \{t \succsim_c s, t \leq_c \alpha, s \leq_c \alpha\}}$$

$$\frac{\emptyset \sqcap \emptyset \rightsquigarrow \emptyset; \emptyset}{\Gamma_1 \sqcap \Gamma_2 \rightsquigarrow \Gamma; \mathcal{C}} \quad \frac{\Gamma_1 \sqcap \Gamma_2 \rightsquigarrow \Gamma; \mathcal{C}}{(\Gamma_1, u : t) \sqcap (\Gamma_2, u : s) \rightsquigarrow \Gamma, u : t; \mathcal{C} \cup \{t =_c s\}}$$

The relation $\Gamma_1 \sqcup \Gamma_2 \rightsquigarrow \Gamma; \mathcal{C}$ combines the type environments Γ_1 and Γ_2 into Γ when the names in $\text{dom}(\Gamma_1) \cup \text{dom}(\Gamma_2)$ are used *both* as specified in Γ_1 *and also* in Γ_2 , so \sqcup is analogous to $+$ in (2.3). When Γ_1 and Γ_2 have disjoint domains, their combination is just their union and no constraints are generated. Any name u that occurs in

Table 3. Constraint generation for expressions and processes.

Expressions		
$\frac{[I-INT]}{n : \mathbf{int} \triangleright \emptyset; \emptyset}$	$\frac{[I-NAME]}{u : \alpha \triangleright u : \alpha; \emptyset}$	
$\frac{[I-PAIR]}{e_i : t_i \triangleright \Gamma_i; \mathcal{C}_i \quad (i=1,2) \quad \Gamma_1 \sqcup \Gamma_2 \rightsquigarrow \Gamma; \mathcal{C}_3}{(e_1, e_2) : t_1 \times t_2 \triangleright \Gamma; \mathcal{C}_1 \cup \mathcal{C}_2 \cup \mathcal{C}_3}$	$\frac{[I-INL]}{e : t \triangleright \Gamma; \mathcal{C}}{\mathbf{inl} \ e : t \oplus \alpha \triangleright \Gamma; \mathcal{C}}$	$\frac{[I-INR]}{e : t \triangleright \Gamma; \mathcal{C}}{\mathbf{inr} \ e : \alpha \oplus t \triangleright \Gamma; \mathcal{C}}$
Processes		
$\frac{[I-IDLE]}{\mathbf{0} \triangleright \emptyset; \emptyset} \quad \frac{[I-IN]}{P \triangleright \Gamma, x : t; \mathcal{C} \quad \Gamma \sqcup u : [t]^{\rho, 0} \rightsquigarrow \Gamma'; \mathcal{C}'}{u(x).P \triangleright \Gamma'; \mathcal{C} \cup \mathcal{C}' \cup \{0 <_{\mathbf{c}} \rho\}}$	$\frac{[I-OUT]}{e : t \triangleright \Gamma; \mathcal{C} \quad \Gamma \sqcup u : [t]^{0, \rho} \rightsquigarrow \Gamma'; \mathcal{C}'}{\bar{u}(e) \triangleright \Gamma'; \mathcal{C}' \cup \{0 <_{\mathbf{c}} \rho\}}$	
$\frac{[I-PAR]}{P_i \triangleright \Gamma_i; \mathcal{C}_i \quad (i=1,2) \quad \Gamma_1 \sqcup \Gamma_2 \rightsquigarrow \Gamma; \mathcal{C}_3}{P_1 P_2 \triangleright \Gamma; \mathcal{C}_1 \cup \mathcal{C}_2 \cup \mathcal{C}_3}$	$\frac{[I-REP]}{P \triangleright \Gamma; \mathcal{C} \quad \Gamma \sqcup \Gamma \rightsquigarrow \Gamma'; \mathcal{C}'}{*P \triangleright \Gamma'; \mathcal{C} \cup \mathcal{C}'}$	$\frac{[I-WEAK]}{P \triangleright \Gamma; \mathcal{C}}{P \triangleright \Gamma, u : \alpha; \mathcal{C} \cup \{\alpha \succ_{\mathbf{c}} \alpha\}}$
$\frac{[I-NEW]}{P \triangleright \Gamma, a : t; \mathcal{C}}{(\mathbf{va})P \triangleright \Gamma; \mathcal{C} \cup \{t =_{\mathbf{c}} [\alpha]^{\rho, \rho}\}}$	$\frac{[I-LET]}{e : t \triangleright \Gamma_1; \mathcal{C}_1 \quad P \triangleright \Gamma_2, x : t_1, y : t_2; \mathcal{C}_2 \quad \Gamma_1 \sqcup \Gamma_2 \rightsquigarrow \Gamma; \mathcal{C}_3}{\mathbf{let} \ x, y = e \ \mathbf{in} \ P \triangleright \Gamma; \mathcal{C}_1 \cup \mathcal{C}_2 \cup \mathcal{C}_3 \cup \{t =_{\mathbf{c}} t_1 \times t_2\}}$	
$\frac{[I-CASE]}{e : t \triangleright \Gamma_1; \mathcal{C}_1 \quad P_i \triangleright \Gamma_i, x_i : t_i; \mathcal{C}_i \quad (i=\mathbf{inl}, \mathbf{inr}) \quad \Gamma_{\mathbf{inl}} \sqcap \Gamma_{\mathbf{inr}} \rightsquigarrow \Gamma_2; \mathcal{C}_2 \quad \Gamma_1 \sqcup \Gamma_2 \rightsquigarrow \Gamma_3; \mathcal{C}_3}{\mathbf{case} \ e \ \mathbf{of} \ \{i \ x_i \Rightarrow P_i\}_{i=\mathbf{inl}, \mathbf{inr}} \triangleright \Gamma_3; \mathcal{C}_1 \cup \mathcal{C}_2 \cup \mathcal{C}_3 \cup \mathcal{C}_{\mathbf{inl}} \cup \mathcal{C}_{\mathbf{inr}} \cup \{t =_{\mathbf{c}} t_{\mathbf{inl}} \oplus t_{\mathbf{inr}}\}}$		

$\text{dom}(\Gamma_1) \cap \text{dom}(\Gamma_2)$ must be used according to compatible types $\Gamma_1(u) \succ \Gamma_2(u)$ and its type must be an upper bound of both $\Gamma_1(u)$ and $\Gamma_2(u)$. In general $\Gamma_1(u)$ and $\Gamma_2(u)$ are type expressions with free type variables, hence these relations cannot be checked right away. Rather, they are symbolically recorded in the set of constraints \mathcal{C} . Note in particular that the combined type of u is unknown and is represented by a fresh type variable that is an upper bound of $\Gamma_1(u)$ and $\Gamma_2(u)$. The relation $\Gamma_1 \sqcap \Gamma_2 \rightsquigarrow \Gamma; \mathcal{C}$ merges the type environments Γ_1 and Γ_2 into Γ when the names in $\text{dom}(\Gamma_1) \cup \text{dom}(\Gamma_2)$ are used in alternative branches of a **case** construct. Note that $\Gamma_1 \sqcap \Gamma_2 \rightsquigarrow \Gamma; \mathcal{C}$ holds if and only if $\Gamma_1(u) = \Gamma_2(u)$ for every $u \in \text{dom}(\Gamma_1) = \text{dom}(\Gamma_2)$. This corresponds to the fact that in [T-CASE] we use the *same* type environment Γ' for typing the two branches of the **case**.

The rules to reconstruct type environments and to generate constraints are presented in Table 3 and derive judgements of the form $e : t \triangleright \Gamma; \mathcal{C}$ for expressions and $P \triangleright \Gamma; \mathcal{C}$ for processes. They closely correspond to those in Table 2; for this and space reasons we will not describe them in detail. In general, unknown uses and types become fresh use and type variables (all variables introduced by the rules are assumed to be fresh), every application of $+$ in Table 2 becomes an application of \sqcup in Table 3, and every assumption on uses and types becomes a constraint. Constraints accumulate from the premises to the conclusion of each rule. In rules [I-INL] and [I-INR] the type of the

Table 4. Constraint solver algorithm.

<p>Input: a set of constraints \mathcal{C}.</p> <p>Output: either fail or a solution of \mathcal{C}.</p> <ol style="list-style-type: none"> 1. Compute $\overline{\mathcal{C}}$; 2. Compute a minimal solution σ_{use} for the use constraints in $\overline{\mathcal{C}}$, or fail if there is none; 3. If $t \sim_c s \in \overline{\mathcal{C}}$ and t, s are proper and have different topmost type constructors, then fail; 4. Let $\sigma_{type} = \{\alpha \mapsto \sup_{\overline{\mathcal{C}}, \sigma_{use}}(\{\alpha\}) \mid \alpha \in \text{dom}(\overline{\mathcal{C}})\}$; 5. Return $\sigma_{use} \cup \sigma_{type}$.

Within each **let** the variable x is assigned a distinct type variable α_i . Eventually, [I-PAR] finds out that x occurs twice, so it records in \mathcal{C}_7 the fact that the two types α_1 and α_2 must be compatible and that the overall type α of x must be an upper bound of both. ■

4 Constraint Solving

In this section we define an algorithm that determines whether a given set of constraints \mathcal{C} is satisfiable and, if this is the case, computes a solution of \mathcal{C} . The algorithm, sketched in Table 4, comprises 5 steps that can be roughly grouped in three phases: *saturation*, *verification*, and *synthesis*. The phases are detailed in the rest of the section.

Saturation (step 1). The \leq_c and \succ_c constraints determined during constraint generation relate type expressions, but they are meant to affect the use variables occurring in these type expressions (recall from (2.2) that every relation \mathcal{R}_{type} between types is the extension of \mathcal{R}_{use} between uses). In order to find all constraints that must hold between use expressions, we *saturate* the set \mathcal{C} with all the constraints that are entailed by those already in \mathcal{C} . Entailment is expressed through a binary relation \models defined as follows:

[E-REFL]	$\{t \mathcal{R}_c s\} \models \{t \mathcal{R}_c t, s \mathcal{R}_c s\}$	$\mathcal{R} \in \{\leq, \sim\}$
[E-SYMM]	$\{t \mathcal{R}_c s\} \models \{s \mathcal{R}_c t\}$	$\mathcal{R} \in \{=, \sim\}$
[E-TRANS]	$\{t_1 \mathcal{R}_c t_2, t_2 \mathcal{R}_c t_3\} \models \{t_1 \mathcal{R}_c t_3\}$	$\mathcal{R} \in \{\leq, \sim\}$
[E-COMP 1]	$\{t_1 =_c t_2, t_2 \succ_c t_3\} \models \{t_1 \succ_c t_3\}$	
[E-COMP 2]	$\{t_1 \leq_c t_2, t_2 \succ_c t_3\} \models \{t_1 \succ_c t_3\}$	
[E-OPER]	$\{t_1 \odot t_2 \mathcal{R}_c s_1 \odot s_2\} \models \{t_1 \mathcal{R}_c s_1, t_2 \mathcal{R}_c s_2\}$	
[E-CHANNEL]	$\{[t]^{K_1, K_2} \mathcal{R}_c [s]^{K_3, K_4}\} \models \{t =_c s, K_1 \mathcal{R}_c K_3, K_2 \mathcal{R}_c K_4\}$	
[E-STRUCT]	$\{t \mathcal{R}_c s\} \models \{t \sim_c s\}$	

The first three rules [E-REFL], [E-SYMM], and [E-TRANS] compute the reflexive, symmetric, and transitive closures of those relations that enjoy such properties. Rule [E-COMP 1] propagates compatibility constraints across equivalent types, and [E-COMP 2] propagates compatibility constraints “downwards” from a type to a smaller one (indeed, it is the case that $\kappa_1 \leq \kappa_2 \succ \kappa_3$ implies $\kappa_1 \succ \kappa_3$). Rule [E-OPER] propagates constraints between composite types to their components. Rule [E-CHANNEL] propagates constraints from type to use expressions and imposes the equality of message types for related channel

types. Finally, rule [E-STRUCT] generates \sim_c constraints between any pair of related type expressions. This is necessary to make sure that all message type equality constraints are generated by [E-CHANNEL], given that \succ is not transitive. We denote by $\overline{\mathcal{C}}$ the smallest set that includes \mathcal{C} and that is closed by the rules [E-*] above. Observe that every \mathcal{C} generated by the rules in Table 3 is finite, and that the entailment rules [E-*] do not change the domain of the set \mathcal{C} being saturated. Therefore, $\overline{\mathcal{C}}$ is always finite and can be computed in finite time by a simple iterative algorithm that repeatedly applies the entailment rules until no new constraints are discovered. We have:

Proposition 4.1. *\mathcal{C} and $\overline{\mathcal{C}}$ are equivalent.*

Verification (steps 2 and 3). In this phase the algorithm verifies whether $\overline{\mathcal{C}}$ is satisfiable and fails if this is not the case. The key observation is that satisfiability of the type constraints does not depend upon *one particular* solution of the use constraints because the previous phase has computed all possible relations that must hold between use expressions. Therefore, we can independently verify the satisfiability of use and type constraints and fail if any of these checks fails.

Recall that there is a finite number of use constraints, which are relationships between use expressions made of a finite number of use variables ranging over a finite domain $\{0, 1, \omega\}$. Therefore, there exists a complete (albeit combinatorial) verification algorithm that determines whether or not the use constraints in $\overline{\mathcal{C}}$ are satisfiable. It is also possible to define an “optimal” algorithm that aims at maximising the number of use variables that are assigned value 1 as opposed to ω . We do not discuss the issues related to solving use constraints any further.

If the use constraints in $\overline{\mathcal{C}}$ are satisfiable, then satisfiability of the type constraints is granted provided that there are no constraints relating types built with different constructors. For example, $\mathbf{int} \leq_c [\alpha]^{k_1, k_2}$ is clearly unsatisfiable. Because of [E-STRUCT] and [E-TRANS], for any pair of types that must be related there is a constraint $t \sim_c s$ in $\overline{\mathcal{C}}$. Therefore, if there is any such constraint where t and s are not type variables and are built using different topmost constructors, then $\overline{\mathcal{C}}$ is for sure unsatisfiable and the algorithm fails.

Proposition 4.2. *If the algorithm fails in this phase, then $\overline{\mathcal{C}}$ is not satisfiable.*

Synthesis (steps 4 and 5). The last phase computes a solution σ_{type} for the type constraints in $\overline{\mathcal{C}}$ given any minimal solution σ_{use} for the use constraints in $\overline{\mathcal{C}}$ determined at step 2 of the algorithm. To compute σ_{type} we need the definitions below:

$$\begin{aligned} \text{cls}_{\mathcal{R}, \mathcal{C}}(T) &\stackrel{\text{def}}{=} \{s \mid s \mathcal{R}_c t \in \mathcal{C} \text{ for some } t \in T \text{ and } s \text{ is proper}\} \\ \text{sup}_{\mathcal{C}, \sigma}(T) &\stackrel{\text{def}}{=} \begin{cases} [\text{sup}_{\mathcal{C}, \sigma}(\{t_i\}_{i \in I})]^{\forall i \in I \sigma \kappa_i, \forall i \in I \sigma \kappa'_i} & \text{if } \text{cls}_{\leq, \mathcal{C}}(T) = \{[t_i]^{k_i, k'_i}\}_{i \in I} \neq \emptyset \\ \text{sup}_{\mathcal{C}, \sigma}(\{t_i\}_{i \in I}) \odot \text{sup}_{\mathcal{C}, \sigma}(\{s_i\}_{i \in I}) & \text{if } \text{cls}_{\leq, \mathcal{C}}(T) = \{t_i \odot s_i\}_{i \in I} \neq \emptyset \\ \text{zero}_{\mathcal{C}, \sigma}(T) & \text{otherwise} \end{cases} \\ \text{zero}_{\mathcal{C}, \sigma}(T) &\stackrel{\text{def}}{=} \begin{cases} [\text{sup}_{\mathcal{C}, \sigma}(\{t_i\}_{i \in I})]^{0,0} & \text{if } \text{cls}_{\sim, \mathcal{C}}(T) = \{[t_i]^{k_i, k'_i}\}_{i \in I} \neq \emptyset \\ \text{zero}_{\mathcal{C}, \sigma}(\{t_i\}_{i \in I}) \odot \text{zero}_{\mathcal{C}, \sigma}(\{s_i\}_{i \in I}) & \text{if } \text{cls}_{\sim, \mathcal{C}}(T) = \{t_i \odot s_i\}_{i \in I} \neq \emptyset \\ \mathbf{int} & \text{otherwise} \end{cases} \end{aligned}$$

The set $\text{cls}_{\mathcal{R}, \mathcal{C}}(T)$ is made of the proper type expressions s such that $s \mathcal{R}_c t \in \mathcal{C}$ for some $t \in T$. Note that not all \mathcal{R} 's are symmetric and that s is the type expression on the left hand side of \mathcal{R}_c . So, $\text{cls}_{\leq, \mathcal{C}}(T)$ is the set of type expressions that are lower bounds of some $t \in T$, while $\text{cls}_{\sim, \mathcal{C}}(T)$ is the set of type expressions that share the same topmost type constructor with some $t \in T$ but have possibly different uses. Note also that $\text{cls}_{\leq, \mathcal{C}}(T) \subseteq \text{cls}_{\sim, \mathcal{C}}(T)$ because $\leq \subseteq \sim$. The algorithm (Table 4) resolves each variable α to $\text{sup}_{\mathcal{C}, \sigma_{use}}(\{\alpha\})$ where, $\text{sup}_{\mathcal{C}, \sigma}(T)$ is, roughly speaking, the least upper bound of the types in T (even though the algorithm always invokes $\text{sup}_{\mathcal{C}, \sigma}$ with a singleton, in general we need to define $\text{sup}_{\mathcal{C}, \sigma}$ over a set of type expressions that are known to be equivalent). There are three cases that determine $\text{sup}_{\mathcal{C}, \sigma}(T)$: if there exists any lower bound for some of the types in T and these lower bounds are either channel or composite types, then $\text{sup}_{\mathcal{C}, \sigma}(T)$ is defined as the least upper bound of such lower bounds (first two cases in the definition of $\text{sup}_{\mathcal{C}, \sigma}$); if there is no lower bound but there exists at least one \sim_c constraint involving any of the types in T , then $\text{sup}_{\mathcal{C}, \sigma}(T)$ is defined as a type that is structurally coherent with such constraints but has use 0 for all of its topmost channel types (third case in the definition of $\text{sup}_{\mathcal{C}, \sigma}$ and first two cases in the definition of $\text{zero}_{\mathcal{C}, \sigma}$); if there are no \sim_c constraints involving any of the types in T , or if some of the types in T have been determined to be structurally coherent with **int**, then $\text{zero}_{\mathcal{C}, \sigma}(T)$ is defined to be **int** (third case in the definition of $\text{zero}_{\mathcal{C}, \sigma}$).

Interpreting $\text{sup}_{\mathcal{C}, \sigma}$ and $\text{zero}_{\mathcal{C}, \sigma}$ as functions is appropriate for presentation (and implementation) purposes, but formally tricky for two reasons: **(1)** the equations given above are mutually dependent and **(2)** they are undefined for some particular T 's (for instance, for $T = \{\mathbf{int}, [\mathbf{int}]^{0,0}\}$ which contains two types with incompatible structures or for $T = \{[\mathbf{int}]^{\omega,0}, [\mathbf{int}]^{1,0}\}$ which contains two types with incompatible uses). Concerning **(1)**, the formal interpretation of the equations above is as a set $\{\alpha_i = t_i\}$ where each α_i has the form $\text{sup}_{\mathcal{C}, \sigma}(T)$ or $\text{zero}_{\mathcal{C}, \sigma}(T)$, $T \subseteq \text{dom}(\mathcal{C})$, and t_i is determined by the right hand side of the equation. We know that this set is always finite because $\text{dom}(\mathcal{C})$ is finite and so is its powerset. Furthermore, $\text{zero}_{\mathcal{C}, \sigma}$ always yields a proper type when it is defined and so does $\text{sup}_{\mathcal{C}, \sigma}$ when it is not defined in terms of $\text{zero}_{\mathcal{C}, \sigma}$. Therefore, the equations in $\{\alpha_i = t_i\}$ are contractive in the sense that there is no infinite chain of equations involving type variables only. In this case, it is known [1] that these equations can be folded into a possibly recursive contractive term using μ 's. Concerning **(2)**, it turns out that, when σ is a solution of the use constraints in \mathcal{C} , $\text{sup}_{\mathcal{C}, \sigma}(T)$ is defined if T is \mathcal{C} -composable and that $\text{zero}_{\mathcal{C}, \sigma}(T)$ is defined if T is \mathcal{C} -compatible, where:

- T is \mathcal{C} -composable if $t \leq_c s \in \mathcal{C}$ or $s \leq_c t \in \mathcal{C}$ or $t \succ_c s \in \mathcal{C}$ for every $t, s \in T$;
- T is \mathcal{C} -compatible if $t \sim_c s \in \mathcal{C}$ for every $t, s \in T$.

Indeed observe that: $\text{cls}_{\leq, \mathcal{C}}(T)$ is \mathcal{C} -composable if so is T by [E-COMP 2]; $\text{cls}_{\sim, \mathcal{C}}(T)$ is \mathcal{C} -compatible if T is \mathcal{C} -composable by [E-STRUCT]; if $\{[t_i]^{k_i, k'_i}\}_{i \in I}$ is \mathcal{C} -composable then the least upper bounds $\bigvee_{i \in I} \sigma \kappa_i$ and $\bigvee_{i \in I} \sigma \kappa'_i$ are defined (consequence of the use constraints generated by [E-CHANNEL] and the hypothesis that σ is a solution of them); if $\{[t_i]^{k_i, k'_i}\}_{i \in I}$ is \mathcal{C} -compatible, then $\{t_i\}_{i \in I}$ is \mathcal{C} -composable (consequence of the $=_c$ constraints generated by [E-CHANNEL]); if $\{t_i \odot s_i\}_{i \in I}$ is \mathcal{C} -composable/compatible, then so are the sets $\{t_i\}_{i \in I}$ and $\{s_i\}_{i \in I}$ by [E-OPER]; the type expressions in \mathcal{C} -composable/compatible sets are built using the same topmost constructor (check in step 3 of the algo-

rithm). Finally, observe that a singleton $\{\alpha\}$ is always \mathcal{C} -composable, so the invocations of $\text{sup}_{\overline{\mathcal{C}}, \sigma_{use}}$ in Table 4 regard well-defined equations. We conclude:

Theorem 4.1. *If the algorithm returns σ , then σ is a minimal solution for \mathcal{C} .*

Each step of the algorithm terminates and if the algorithm fails it is because \mathcal{C} has no solution (Proposition 4.2). Therefore:

Corollary 4.1 (completeness). *If \mathcal{C} is satisfiable, the algorithm returns a solution.*

Example 4.1. The saturation of the constraint set \mathcal{C} computed in Example 3.1 contains, among others, the constraints $[\mathbf{int}]^{0, \rho_1} \times \gamma \succ_c \beta \times [\mathbf{int}]^{0, \rho_2}$ and consequently $[\mathbf{int}]^{0, \rho_1} \succ_c \beta$ and $\gamma \succ_c [\mathbf{int}]^{0, \rho_2}$ by [E-COMP 1] and [E-COMP 2]. An optimal solution of the use constraints in $\overline{\mathcal{C}}$ is $\sigma_{use} \stackrel{\text{def}}{=} \{\rho_1 \mapsto 1, \rho_2 \mapsto 1, \rho_3 \mapsto 1\}$. From this we obtain

$$\text{sup}_{\overline{\mathcal{C}}, \sigma_{use}}(\{\alpha\}) = [\mathbf{int}]^{0,1} \times [\mathbf{int}]^{0,1}$$

indicating that the pair of channels received from a is shared by the two `let` processes in such a way that each of the two channels contained therein is used exactly once. ■

5 Concluding Remarks

Previous works on the linear π -calculus either ignore composite types [11,7] or are based on an interpretation of linearity that limits data sharing and parallelism [5,6]. Recursive types have also been neglected, despite their prominent role for describing complex interactions occurring on linear channels [2]. In this work we extend the linear π -calculus with both composite and recursive types and we adopt a more relaxed attitude towards linearity that fosters data sharing and parallelism while preserving complete type reconstruction. The extension is a very natural one, as witnessed by the fact that our type system uses essentially the same rules of previous works, the main novelty being a different type composition operator. This small change has nonetheless non-trivial consequences on the reconstruction algorithm, which must reconcile the propagation of constraints across composite types with the impossibility to rely on plain type unification due to the fact that different occurrences of the same identifier may be assigned different types and because of recursive types. Technically, we tackle these problems by expressing type combination in previous works (which is a ternary relation $t_1 + t_2 = t_3$) in terms of two simpler binary relations, namely compatibility $t_1 \asymp t_2$ and order $t_i \leq t_3$, and by seeking for *minimal* solutions of the constraint set. Our extension also gives renewed relevance to types like $[t]^{0,0}$. In previous works these types were admitted but essentially useless: channels with such types could only be passed around in messages without actually ever being used. That is, they could be erased without affecting processes. In our type system, it is the existence of these types that enables the sharing of structured data (see the decomposition of t_{list} into t_{even} and t_{odd} in Section 1).

Given that sessions [3,4] can be fully encoded into the linear π -calculus [9,2], we also indirectly provide a complete reconstruction algorithm for equi-recursive session types without subtyping. Interestingly, the concept of *duality*, which is a source of significant complication of the type reconstruction algorithm for *finite* session types [12], is simplified by the encoding, where it reduces to compatibility.

To assess the feasibility of the approach, we have developed a prototype based on the naïve constraint saturation and combinatorial verification of use expressions described in Section 4. Regarding the complexity of the reconstruction algorithm, the most critical aspects are the solution of use constraints and the computation of the transitive closure of type constraints. While polynomial algorithms are known for the latter, the former problem entails, in principle, an exponential cost. Preliminary experiments with the prototype implementation of the algorithm have shown that, in both cases, constraints can often be partitioned into relatively small independent subsets that can be solved in isolation (the prototype already supports such partitioning for use constraints). This property paves the way for significant performance improvements.

Acknowledgements. The author is grateful to the FoSSaCS'14 reviewers for their detailed and insightful comments. This work has been partially supported by ICT COST Action IC1201 BETTY, MIUR project CINA, and Ateneo/CSP project SALT.

References

1. B. Courcelle. Fundamental properties of infinite trees. *Theor. Comp. Sci.*, 25:95–169, 1983.
2. O. Dardha, E. Giachino, and D. Sangiorgi. Session types revisited. In *PPDP'12*, pages 139–150. ACM, 2012.
3. K. Honda. Types for dyadic interaction. In *CONCUR'93*, LNCS 715, pages 509–523. Springer, 1993.
4. K. Honda, V. T. Vasconcelos, and M. Kubo. Language primitives and type disciplines for structured communication-based programming. In *ESOP'98*, LNCS 1381, pages 122–138. Springer, 1998.
5. A. Igarashi. Type-based analysis of usage of values for concurrent programming languages, 1997. Available at <http://www.sato.kuis.kyoto-u.ac.jp/~igarashi/papers/>.
6. A. Igarashi and N. Kobayashi. Type-based analysis of communication for concurrent programming languages. In *SAS'97*, LNCS 1302, pages 187–201. Springer, 1997.
7. A. Igarashi and N. Kobayashi. Type Reconstruction for Linear π -Calculus with I/O Subtyping. *Inf. and Comp.*, 161(1):1–44, 2000.
8. N. Kobayashi. Quasi-linear types. In *POPL'99*, pages 29–42. ACM, 1999.
9. N. Kobayashi. Type systems for concurrent programs. In *10th Anniversary Colloquium of UNU/IIST*, LNCS 2757, pages 439–453. Springer, 2002. Extended version at <http://www.kb.ecei.tohoku.ac.jp/~koba/papers/tutorial-type-extended.pdf>.
10. N. Kobayashi. A new type system for deadlock-free processes. In *CONCUR'06*, LNCS 4137, pages 233–247. Springer, 2006.
11. N. Kobayashi, B. C. Pierce, and D. N. Turner. Linearity and the pi-calculus. *ACM Trans. Program. Lang. Syst.*, 21(5):914–947, 1999.
12. L. G. Mezzina. How to infer finite session types in a calculus of services and sessions. In *COORDINATION'08*, LNCS 5052, pages 216–231. Springer, 2008.
13. U. Nestmann and M. Steffen. Typing confluence. In *FMICS'97*, pages 77–101, 1997. Also available as report ERCIM-10/97-R052, European Research Consortium for Informatics and Mathematics, 1997.
14. D. Sangiorgi and D. Walker. *The Pi-Calculus - A theory of mobile processes*. Cambridge University Press, 2001.
15. D. N. Turner, P. Wadler, and C. Mossin. Once upon a type. In *FPCA'95*, pages 1–11, 1995.

A Supplement to Section 2

In order to prove Theorem A.1 we need a series of standard auxiliary results, including a weakening lemma (Lemma A.1), type preservation under structural congruence (Lemma A.2), and substitution (Lemma A.3).

Lemma A.1. *If $\Gamma \vdash P$ and $\Gamma' \asymp \Gamma$ and $\Gamma + \Gamma'$ is defined, then $\Gamma + \Gamma' \vdash P$.*

Proof. This is a standard induction on the derivation of $\Gamma \vdash P$ where we assume, without loss of generality, that $\text{bn}(P) \cap \text{dom}(\Gamma) = \emptyset$ (recall that we identify processes modulo renaming of bound names). \square

Lemma A.2. *If $\Gamma \vdash P$ and $P \equiv Q$, then $\Gamma \vdash Q$.*

Proof. We show two interesting cases only.

$\boxed{\text{[S-PAR 1]}}$ Then $P = \mathbf{0} \mid Q$. From [T-PAR] we deduce $\Gamma = \Gamma_1 + \Gamma_2$ and $\Gamma_1 \vdash \mathbf{0}$ and $\Gamma_2 \vdash Q$. From [T-IDLE] we deduce $\Gamma_1 \asymp \Gamma_1$. We conclude by Lemma A.1.

$\boxed{\text{[S-REP]}}$ Then $P = *P' \equiv *P' \mid P' = Q$. From the hypothesis $\Gamma \vdash P$ we deduce $\Gamma = \Gamma' + \Gamma'$ for some Γ' such that $\Gamma' \vdash P'$. We conclude with an application of [T-PAR] and observing that $\Gamma = \Gamma' + \Gamma' = \Gamma' + \Gamma' + \Gamma' = \Gamma + \Gamma'$. \square

Lemma A.3 (substitution). *If $\Gamma, x : t \vdash P$ and $\Gamma' \vdash v : t$ and $\Gamma + \Gamma'$ is defined, then $\Gamma + \Gamma' \vdash P\{v/x\}$.*

Proof. The proof is unremarkable, except for the following property of the type system: $t \asymp t$ implies $\Gamma' \asymp \Gamma'$, which can be easily proved by induction on the derivation of $\Gamma' \vdash v : t$. \square

In order to state the subject reduction result we must take into account the fact that the type of a linear channel changes whenever a synchronization occurs on it, because the two 1 uses are reset to 0, indicating that the channel can no longer be used. To this aim we define a reduction relation over type environments. In particular, let $\xrightarrow{\ell}$ be the least relation between type environments such that

$$\Gamma, a : [t]^{1,1} \xrightarrow{a} \Gamma, a : [t]^{0,0} \quad \text{and} \quad \Gamma, a : [t]^{\omega,\omega} \xrightarrow{a} \Gamma, a : [t]^{\omega,\omega} \quad \text{and} \quad \Gamma \xrightarrow{\tau} \Gamma$$

We say that Γ is **balanced** if every type $[t]^{\kappa_1, \kappa_2}$ in the range of Γ where $0 < \kappa_1$ and $0 < \kappa_2$ is such that $\kappa_1 = \kappa_2$.

Lemma A.4. *The following properties hold:*

1. *If Γ is balanced and $\Gamma \xrightarrow{\ell} \Gamma'$, then Γ' is balanced.*
2. *If $\Gamma \xrightarrow{\ell} \Gamma'$ and $\Gamma + \Gamma''$ is defined, then $\Gamma + \Gamma'' \xrightarrow{\ell} \Gamma' + \Gamma''$.*

Proof. Easy consequences of the definition of $\xrightarrow{\ell}$ over type environments. \square

Theorem A.1 (subject reduction). *Let $\Gamma \vdash P$ and Γ balanced and $P \xrightarrow{\ell} Q$. Then $\Gamma' \vdash Q$ for some Γ' such that $\Gamma \xrightarrow{\ell} \Gamma'$.*

Proof. By induction on the derivation of $P \xrightarrow{\ell} Q$ and by cases on the last rule applied. We only show a few interesting cases; the others are either similar or simpler.

[R-COMM] Then $P = \bar{a}\langle v \rangle | a(x).R$ and $\ell = a$. From [T-PAR] we deduce $\Gamma = \Gamma_1 + \Gamma_2$ where $\Gamma_1 \vdash \bar{a}\langle v \rangle$ and $\Gamma_2 \vdash a(x).R$. From [T-OUT] we deduce $\Gamma_1 = \Gamma'_1 + a : [t]^{0,\kappa}$ and $0 < \kappa$ and $\Gamma'_1 \vdash v : t$. From [T-IN] we deduce $\Gamma_2 = \Gamma'_2 + a : [t]^{\kappa,0}$ and $\Gamma'_2, x : t \vdash R$. The reason why we know that we have the same t and κ in Γ_1 and Γ_2 comes from the definition of $+$ and the fact that these two environments were combined together in Γ which is balanced. Let

$$\Gamma_0 \stackrel{\text{def}}{=} \begin{cases} a : [t]^{0,0} & \text{if } \kappa = 1 \\ a : [t]^{\omega,\omega} & \text{if } \kappa = \omega \end{cases} \quad \text{and} \quad \Gamma' \stackrel{\text{def}}{=} \Gamma_0 + \Gamma'_1 + \Gamma'_2$$

and observe that $\Gamma_0 \asymp \Gamma_0$, that $\Gamma_0 + \Gamma'_1 + \Gamma'_2$ is defined and $\Gamma \xrightarrow{a} \Gamma'$. By Lemma A.3 and Lemma A.1 we conclude $\Gamma' \vdash R\{v/x\}$.

[R-LET] Then $P = \mathbf{let} \ x, y = v, w \ \mathbf{in} \ R \xrightarrow{\tau} R\{v, w/x, y\} = Q$. From [T-LET] we deduce $\Gamma = \Gamma_1 + \Gamma_2$ and $\Gamma_1 \vdash v, w : t \times s$ and $\Gamma_2, x : t, y : s \vdash R$. From [T-PAIR] we deduce $\Gamma_1 = \Gamma_{11} + \Gamma_{12}$ and $\Gamma_{11} \vdash v : t$ and $\Gamma_{12} \vdash w : s$. We conclude $\Gamma \vdash Q$ by Lemma A.3.

[R-PAR] Then $P = P_1 | P_2$ and $P_1 \xrightarrow{\ell} P'_1$ and $Q = P'_1 | P_2$. From [T-PAR] we deduce $\Gamma = \Gamma_1 + \Gamma_2$ and $\Gamma_i \vdash P_i$ where Γ_i is balanced for $i \in \{1, 2\}$. By induction hypothesis we deduce $\Gamma'_1 \vdash P'_1$ for some Γ'_1 such that $\Gamma_1 \xrightarrow{\ell} \Gamma'_1$. By Proposition A.4 we deduce that $\Gamma \xrightarrow{\ell} \Gamma'_1 + \Gamma_2$. We conclude $\Gamma' \vdash Q$ by taking $\Gamma' = \Gamma'_1 + \Gamma_2$. \square

Note that Theorem A.1 establishes not only a subject reduction result, but also a soundness result because it says that a channel is used no more than what is allowed by its type. It is also possible to establish some basic safety properties, in particular that well-typed **let**'s and **case**'s without free variables always reduce. The proof is easy, we omit the details.

B Supplement to Section 3

Lemma B.1. *If $\Gamma_1 \sqcup \Gamma_2 \rightsquigarrow \Gamma; \mathcal{C}$ and σ is a minimal solution for \mathcal{C} , then $\Gamma\sigma = \Gamma_1\sigma + \Gamma_2\sigma$.*

Proof. By induction on the derivation of $\Gamma_1 \sqcup \Gamma_2 \rightsquigarrow \Gamma; \mathcal{C}$.

dom(Γ_1) \cap **dom**(Γ_2) = \emptyset Then $\Gamma = \Gamma_1, \Gamma_2$ and we conclude $\Gamma\sigma = (\Gamma_1, \Gamma_2)\sigma = \Gamma_1\sigma, \Gamma_2\sigma = \Gamma_1\sigma + \Gamma_2\sigma$.

$\Gamma_1 = \Gamma'_1, u : t$ and $\Gamma_2 = \Gamma'_2, u : s$ Then $\Gamma'_1 \sqcup \Gamma'_2 \rightsquigarrow \Gamma'; \mathcal{C}'$ and $\Gamma = \Gamma', u : \alpha$ and $\mathcal{C} = \mathcal{C}' \cup \{t \succ_c s, t \leq_c \alpha, s \leq_c \alpha\}$ where α is fresh. Since σ is a minimal solution for \mathcal{C} , we deduce $\sigma(\alpha) = t\sigma \vee s\sigma$. Also, σ is a minimal solution for \mathcal{C}' because α is fresh, hence it does not occur in \mathcal{C}' . By induction hypothesis we deduce $\Gamma'\sigma = \Gamma'_1\sigma + \Gamma'_2\sigma$. We conclude $\Gamma\sigma = (\Gamma', u : \alpha)\sigma = \Gamma'\sigma, u : \sigma(\alpha) = (\Gamma'_1\sigma + \Gamma'_2\sigma), u : t\sigma \vee s\sigma = \Gamma_1\sigma + \Gamma_2\sigma$. \square

Lemma B.2. *If $\Gamma_1 \sqcap \Gamma_2 \rightsquigarrow \Gamma; \mathcal{C}$ and σ is a solution for \mathcal{C} , then $\Gamma\sigma = \Gamma_1\sigma = \Gamma_2\sigma$.*

Proof. Straightforward consequence of the definition of $\Gamma_1 \sqcap \Gamma_2 \rightsquigarrow \Gamma; \mathcal{C}$. \square

Lemma B.3. *If $e : t \triangleright \Gamma; \mathcal{C}$ and σ is a minimal solution of \mathcal{C} , then $\Gamma \sigma \vdash e : t \sigma$.*

Proof. By induction on the derivation of $e : t \triangleright \Gamma; \mathcal{C}$ and by cases on the last rule applied. We only show two significant cases.

[I-NAME] Then $e = u$ and $t = \alpha$ fresh and $\Gamma = u : \alpha$ and $\mathcal{C} = \emptyset$. We have $\Gamma \sigma = u : \sigma(\alpha)$ and $t \sigma = \sigma(\alpha)$, hence we conclude $\Gamma \sigma \vdash e : t \sigma$.

[I-PAIR] Then $e = e_1, e_2$ and $t = t_1 \times t_2$ and $\mathcal{C} = \mathcal{C}_1 \cup \mathcal{C}_2 \cup \mathcal{C}_3$ where $\Gamma_1 \sqcap \Gamma_2 \rightsquigarrow \Gamma; \mathcal{C}_3$ and $e_i : t_i \triangleright \Gamma_i; \mathcal{C}_i$ for $i = 1, 2$. Since every distinct derivation introduces fresh type variables we know that σ is a minimal solution for both Γ_1 and Γ_2 . By induction hypothesis we deduce $\Gamma_i \sigma \vdash e : t_i \sigma$ for $i = 1, 2$. From Lemma B.1 we obtain $\Gamma \sigma = \Gamma_1 \sigma + \Gamma_2 \sigma$. We conclude with an application of [T-PAIR]. \square

Theorem B.1 (Theorem 3.1). *If $P \triangleright \Gamma; \mathcal{C}$ and σ is a minimal solution for \mathcal{C} , then $\Gamma \sigma \vdash P$.*

Proof. By induction on the derivation of $P \triangleright \Gamma; \mathcal{C}$ and by cases on the last rule applied.

[I-IDLE] Then $P = \mathbf{0}$ and $\Gamma = \emptyset$ and $\mathcal{C} = \emptyset$. We conclude with an application of [T-IDLE].

[I-IN] Then $P = u(x).Q$ and $Q \triangleright \Gamma', x : t; \mathcal{C}_1$ and $\Gamma' \sqcup u : [t]^{\rho, 0} \rightsquigarrow \Gamma; \mathcal{C}_2$ and $\mathcal{C} = \mathcal{C}_1 \cup \mathcal{C}_2 \cup \{0 <_c \rho\}$. By induction hypothesis we deduce $\Gamma' \sigma, x : t \sigma \vdash Q$. By Lemma B.1 we deduce $\Gamma \sigma = \Gamma' \sigma + u : [t \sigma]^{\sigma(\rho), 0}$. From the hypothesis that σ is a solution for \mathcal{C}_0 we know $0 < \sigma(\rho)$. We conclude with an application of [T-IN].

[I-OUT] Then $P = \bar{u}(e)$ and $e : t \triangleright \Gamma'; \mathcal{C}_1$ and $\Gamma' \sqcup u : [t]^{0, \rho} \rightsquigarrow \Gamma; \mathcal{C}_2$ and $\mathcal{C} = \mathcal{C}_1 \cup \mathcal{C}_2 \cup \{0 <_c \rho\}$. From Lemma B.3 we deduce $\Gamma' \sigma \vdash e : t \sigma$. From Lemma B.1 we deduce $\Gamma \sigma = \Gamma' \sigma + u : [t \sigma]^{0, \sigma(\rho)}$. From the hypothesis that σ is a solution of \mathcal{C}_0 we know $0 < \sigma(\rho)$. We conclude with an application of [T-OUT].

[I-PAR] Then $P = P_1 \mid P_2$ and $P_i \triangleright \Gamma_i; \mathcal{C}_i$ for $i = 1, 2$ and $\Gamma_1 \sqcap \Gamma_2 \rightsquigarrow \Gamma; \mathcal{C}_3$ and $\mathcal{C} = \mathcal{C}_1 \cup \mathcal{C}_2 \cup \mathcal{C}_3$. By induction hypothesis we deduce $\Gamma_i \sigma \vdash P_i$ for $i = 1, 2$. By Lemma B.1 we deduce $\Gamma \sigma = \Gamma_1 \sigma + \Gamma_2 \sigma$. We conclude with an application of [T-PAR].

[I-REP] Then $P = *Q$ and $Q \triangleright \Gamma'; \mathcal{C}_1$ and $\Gamma' \sqcup \Gamma' \rightsquigarrow \Gamma; \mathcal{C}_2$ and $\mathcal{C} = \mathcal{C}_1 \cup \mathcal{C}_2$. By induction hypothesis we deduce $\Gamma' \sigma \vdash Q$. By Lemma B.1 we deduce $\Gamma \sigma = \Gamma' \sigma + \Gamma' \sigma$. We conclude with an application of [T-REP].

[I-NEW] Then $P = (va)Q$ and $Q \triangleright \Gamma, a : [t]^{\kappa_1, \kappa_2}; \mathcal{C}'$ and $\mathcal{C} = \mathcal{C}' \cup \{\kappa_1 =_c \kappa_2\}$. By induction hypothesis we deduce $\Gamma \sigma, a : [t \sigma]^{\kappa_1 \sigma, \kappa_2 \sigma} \vdash Q$. Since σ is a solution of \mathcal{C}_0 we know that $\kappa_1 \sigma = \kappa_2 \sigma$. We conclude with an application of [T-NEW].

[I-WEAK] Then $\Gamma = \Gamma', u : \alpha$ and $\mathcal{C} = \mathcal{C}' \cup \{\alpha \asymp_c \alpha\}$ where α is fresh and $P \triangleright \Gamma'; \mathcal{C}'$. By induction hypothesis we deduce $\Gamma' \sigma \vdash P$. Since σ is a solution of \mathcal{C}_0 we know that $\sigma(\alpha) \asymp \sigma(\alpha)$. Since $u \notin \text{dom}(\Gamma')$ we know that $\Gamma' \sigma + u : \sigma(\alpha)$ is defined. By Lemma A.1 we conclude $\Gamma' \sigma, u : \sigma(\alpha) \vdash P$.

[I-LET] Then $P = \text{let } x, y = e \text{ in } Q$ and $e : t \triangleright \Gamma_1; \mathcal{C}_1$ and $Q \triangleright \Gamma_2, x : t_1, y : t_2; \mathcal{C}_2$ and $\Gamma_1 \sqcap \Gamma_2 \rightsquigarrow \Gamma; \mathcal{C}_3$ and $\mathcal{C} = \mathcal{C}_1 \cup \mathcal{C}_2 \cup \mathcal{C}_3 \cup \{t =_c t_1 \times t_2\}$. By Lemma B.3 we deduce

$\Gamma_1\sigma \vdash e : t\sigma$. By induction hypothesis we deduce $\Gamma_2\sigma, x : t_1\sigma, y : t_2\sigma \vdash Q$. Since σ is a solution of \mathcal{C} we know that $t\sigma = t_1\sigma \times t_2\sigma$. We conclude with an application of [T-LET].

[I-CASE] Then $P = \text{case } e \text{ of } \{i x_i \Rightarrow P_i\}_{i=\text{inl}, \text{inr}}$ and $e : t \triangleright \Gamma_1; \mathcal{C}_1$ and $P_i \triangleright \Gamma_i, x_i : t_i; \mathcal{C}_i$ for $i = \text{inl}, \text{inr}$ and $\Gamma_{\text{inl}} \sqcap \Gamma_{\text{inr}} \rightsquigarrow \Gamma_2; \mathcal{C}_2$ and $\Gamma_1 \sqcup \Gamma_2 \rightsquigarrow \Gamma; \mathcal{C}_3$ and $\mathcal{C} = \mathcal{C}_1 \cup \mathcal{C}_2 \cup \mathcal{C}_3 \cup \{t =_c t_{\text{inl}} \oplus t_{\text{inr}}\}$. By Lemma B.3 we deduce $\Gamma_1\sigma \vdash e : t\sigma$. By induction hypothesis we deduce $\Gamma_i\sigma \vdash P_i$ for $i = \text{inl}, \text{inr}$. By Lemma B.2 we deduce $\Gamma_{\text{inl}}\sigma = \Gamma_{\text{inr}}\sigma = \Gamma_2\sigma$. By Lemma B.1 we deduce $\Gamma\sigma = \Gamma_1\sigma + \Gamma_2\sigma$. We conclude with an application of [T-CASE]. \square

Lemma B.4. *If $\Gamma_1\sigma + \Gamma_2\sigma$ is defined, then there exist Γ, \mathcal{C} , and $\sigma' \supseteq \sigma$ such that $\Gamma_1 \sqcup \Gamma_2 \rightsquigarrow \Gamma; \mathcal{C}$ and $\Gamma\sigma' = \Gamma_1\sigma' + \Gamma_2\sigma'$ and σ' is solution of \mathcal{C} .*

Proof. By induction on Γ_1 and Γ_2 .

$\text{dom}(\Gamma_1) \cap \text{dom}(\Gamma_2) = \emptyset$ We conclude by taking $\Gamma = \Gamma_1, \Gamma_2$ and $\mathcal{C} = \emptyset$ and $\sigma' = \sigma$ and by observing that $\Gamma_1 \sqcup \Gamma_2 \rightsquigarrow \Gamma; \emptyset$.

$\Gamma_1 = \Gamma'_1, u : t$ and $\Gamma_2 = \Gamma'_2, u : s$ Since $\Gamma_1\sigma + \Gamma_2\sigma$ is defined, we know that $\Gamma'_1\sigma + \Gamma'_2\sigma$ is also defined and furthermore $t\sigma \asymp_c s\sigma$ and $(\Gamma_1\sigma + \Gamma_2\sigma)(u) = t\sigma \vee s\sigma$. By induction hypothesis we deduce that there exist Γ', \mathcal{C}' , and $\sigma'' \supseteq \sigma$ such that $\Gamma'_1 \sqcup \Gamma'_2 \rightsquigarrow \Gamma'; \mathcal{C}'$ and $\Gamma'\sigma'' = \Gamma'_1\sigma'' + \Gamma'_2\sigma''$ and σ'' is a solution of \mathcal{C}' . Take $\Gamma = \Gamma', u : \alpha$ where α is fresh, $\mathcal{C} = \mathcal{C}' \cup \{t \asymp_c s, t \leq_c \alpha, s \leq_c \alpha\}$ and $\sigma' = \sigma'' \cup \{\alpha \mapsto t\sigma \vee s\sigma\}$. We conclude by observing that σ' is a solution of \mathcal{C} . \square

Theorem B.2 (Theorem 3.2). *If $\Gamma \vdash P$, then there exist Γ_0, \mathcal{C}_0 and σ that is a solution of \mathcal{C}_0 such that $P \triangleright \Gamma_0; \mathcal{C}_0$ and $\Gamma = \Gamma_0\sigma$.*

Proof. By induction on the derivation of $\Gamma \vdash P$ and by cases on the last rule applied. We only show a few cases, the others being analogous.

[T-IDLE] Then $P = \mathbf{0}$ and $\Gamma \asymp \Gamma$. Let $\Gamma = \{u_i : t_i\}_{i \in I}$ where $I = \{1, \dots, n\}$. Let $\Gamma_0 = \{u_i : \alpha_i\}_{i \in I}$ and $\mathcal{C}_0 = \{\alpha_i \asymp_c t_i\}_{i \in I}$ and $\sigma = \{\alpha_i \mapsto t_i\}_{i \in I}$ where the α_i 's are all fresh variables. By repeated applications of [I-WEAK] and one application of [I-IDLE] we derive $\mathbf{0} \triangleright \Gamma_0; \mathcal{C}_0$. We conclude by observing that σ is a solution of \mathcal{C}_0 and $\Gamma = \Gamma_0\sigma$.

[T-IN] Then $P = u(x).Q$ and $\Gamma', x : t \vdash Q$ and $\Gamma = \Gamma' + u : [t]^{\kappa, 0}$ and $0 < \kappa$. By induction hypothesis we deduce that there exist $\Gamma'_0, t'_0, \mathcal{C}'_0$, and σ' that is a solution of \mathcal{C}'_0 and $Q \triangleright \Gamma'_0, x : t'_0; \mathcal{C}'_0$ and $\Gamma' = \Gamma'_0\sigma'$ and $t = t'_0\sigma'$. Let $\sigma'' = \sigma' \cup \{\rho \mapsto \kappa\}$ where ρ is fresh. By Lemma B.4 we deduce that there exist Γ_0, \mathcal{C} , and $\sigma \supseteq \sigma''$ such that $\Gamma'_0 \sqcup u : [t'_0]^{\kappa, 0} \rightsquigarrow \Gamma_0; \mathcal{C}$ and $\Gamma_0\sigma = \Gamma'_0\sigma + u : [t'_0\sigma]^{\kappa, 0} = \Gamma$ and σ is a solution of \mathcal{C} . We conclude $P \triangleright \Gamma_0; \mathcal{C}_0$ with an application of [I-IN] by taking $\mathcal{C}_0 = \mathcal{C}'_0 \cup \mathcal{C} \cup \{0 <_c \rho\}$ and by observing that σ is a solution of \mathcal{C}_0 .

[T-PAR] Then $P = P_1 | P_2$ and $\Gamma_i \vdash P_i$ for $i = 1, 2$ and $\Gamma = \Gamma_1 + \Gamma_2$. By induction hypothesis we deduce that, for every $i = 1, 2$, there exist Γ'_i and \mathcal{C}_i and σ_i that is solution of \mathcal{C}_i such that $P_i \triangleright \Gamma'_i; \mathcal{C}_i$ and $\Gamma_i = \Gamma'_i\sigma_i$. We also know that $\text{dom}(\sigma_1) \cap \text{dom}(\sigma_2) = \emptyset$ because type/use variables are always chosen fresh. Let $\sigma' = \sigma_1 \cup \sigma_2$. By Lemma B.4 we deduce that there exist Γ', \mathcal{C}' , and $\sigma \supseteq \sigma'$ such that $\Gamma_1 \sqcup \Gamma_2 \rightsquigarrow \Gamma'; \mathcal{C}'$ and $\Gamma'\sigma = \Gamma_1\sigma + \Gamma_2\sigma$ and

σ is a solution of \mathcal{C}' . We conclude by taking $\mathcal{C} = \mathcal{C}_1 \cup \mathcal{C}_2 \cup \mathcal{C}'$ with an application of [I-PAR].

[T-REP] Then $P = *Q$ and $\Gamma' \vdash Q$ and $\Gamma = \Gamma' + \Gamma'$. By induction hypothesis we deduce that there exist $\Gamma'_0, \mathcal{C}'_0$, and σ' that is a solution of \mathcal{C}'_0 such that $Q \triangleright \Gamma'_0; \mathcal{C}'_0$ and $\Gamma' = \Gamma'_0 \sigma'$. By Lemma B.4 we deduce that there exist $\Gamma_0, \mathcal{C}', \sigma \supseteq \sigma'$ such that $\Gamma'_0 \sqcup \Gamma'_0 \rightsquigarrow \Gamma_0; \mathcal{C}'$ and $\Gamma_0 \sigma = \Gamma'_0 \sigma + \Gamma'_0 \sigma = \Gamma$ and σ is a solution of \mathcal{C}' . We conclude $P \triangleright \Gamma_0; \mathcal{C}$ with an application of [I-REP] by taking $\mathcal{C} = \mathcal{C}'_0 \cup \mathcal{C}'$. \square

C Supplement to Section 4

Definition C.1. We say that \mathcal{C} is **closed for compatibility** if for every $t, s_1, s_2 \in \text{dom}(\overline{\mathcal{C}})$ and $s_i \leq_c t \in \overline{\mathcal{C}}$ for $i = 1, 2$ we have either $s_1 \leq_c s_2$ or $s_2 \leq_c s_1$ or $s_1 \succ_c s_2$.

Lemma C.1. Let \mathcal{C} be a set of constraints generated by the reconstruction system in Table 3. Then \mathcal{C} is closed for compatibility.

Proof. Follows by a simple analysis of the sets of constraints generated by reconstruction system. \square

Lemma C.2. If \mathcal{C} is closed for compatibility, then so is $\overline{\mathcal{C}}$.

Proof. Consequence of the hypothesis that \mathcal{C} is closed for compatibility and the fact that $\overline{\mathcal{C}}$ is closed under [E-COMP 2]. \square

Lemma C.3 (Proposition 4.1). \mathcal{C} and $\overline{\mathcal{C}}$ are equivalent.

Proof. Follows from the saturation laws [E-*] and the fact that each $\mathcal{R}_{\text{type}}$ is the largest relation that satisfies (2.2). \square

Lemma C.4. Let $\emptyset \neq T \subseteq S$ where S is composable. The following properties hold:

1. $\text{cls}_{\sim, \mathcal{C}}(T, \mathcal{C}) = \text{cls}_{\sim, \mathcal{C}}(S, \mathcal{C})$;
2. If $t, s \in S$ implies $t =_c s \in \mathcal{C}$, then $\text{cls}_{\mathcal{R}, \mathcal{C}}(T, \mathcal{C}) = \text{cls}_{\mathcal{R}, \mathcal{C}}(S, \mathcal{C})$ for every $\mathcal{R} \in \{\leq, \sim\}$.

Proof. Easy consequences of the definition of $\text{cls}_{\mathcal{R}, \mathcal{C}}$. \square

Lemma C.5. Let \mathcal{C} be a saturated set of constraints that has passed the checks at steps 2 and 3 of the algorithm in Table 4 and let σ be a solution of the use constraints in \mathcal{C} . Then $\text{sup}_{\mathcal{C}, \sigma}(\{t\}) \mathcal{R} \text{sup}_{\mathcal{C}, \sigma}(\{s\})$ for every $t \mathcal{R}_c s \in \mathcal{C}$.

Proof. In this proof we write $T \mathcal{R}_c S \in \mathcal{C}$ if $t \mathcal{R}_c s \in \mathcal{C}$ for every $t \in T$ and $s \in S$. Let

$$\begin{aligned} \mathcal{C}(\leq) \stackrel{\text{def}}{=} & \{(\text{sup}_{\mathcal{C}, \sigma}(T), \text{sup}_{\mathcal{C}, \sigma}(S)) \mid \emptyset \neq T \subseteq S \subseteq \text{dom}(\mathcal{C}) \text{ and } S \text{ is composable}\} \\ & \cup \{(\text{zero}_{\mathcal{C}, \sigma}(T), \text{sup}_{\mathcal{C}, \sigma}(S)) \mid \emptyset \neq T, S \subseteq \text{dom}(\mathcal{C}) \text{ and } T \sim_c S \text{ and } S \text{ is composable}\} \\ & \cup \{(\text{zero}_{\mathcal{C}, \sigma}(T), \text{zero}_{\mathcal{C}, \sigma}(S)) \mid \emptyset \neq T, S \subseteq \text{dom}(\mathcal{C}) \text{ and } T \sim_c S\} \end{aligned}$$

$$\begin{aligned} \mathcal{C}(\succ) \stackrel{\text{def}}{=} & \{(\text{sup}_{\mathcal{C}, \sigma}(T), \text{sup}_{\mathcal{C}, \sigma}(S)) \mid \emptyset \neq T, S \subseteq \text{dom}(\mathcal{C}) \text{ and } T \succ_c S \in \mathcal{C}\} \\ & \cup \{(\text{zero}_{\mathcal{C}, \sigma}(T), \text{sup}_{\mathcal{C}, \sigma}(S)) \mid \emptyset \neq T, S \subseteq \text{dom}(\mathcal{C}) \text{ and } T \succ_c S \in \mathcal{C}\} \\ & \cup \{(\text{sup}_{\mathcal{C}, \sigma}(T), \text{zero}_{\mathcal{C}, \sigma}(S)) \mid \emptyset \neq T, S \subseteq \text{dom}(\mathcal{C}) \text{ and } T \succ_c S \in \mathcal{C}\} \\ & \cup \{(\text{zero}_{\mathcal{C}, \sigma}(T), \text{zero}_{\mathcal{C}, \sigma}(S)) \mid \emptyset \neq T, S \subseteq \text{dom}(\mathcal{C}) \text{ and } T \succ_c S \in \mathcal{C}\} \end{aligned}$$

and observe that $t \mathcal{R}_c s \in \mathcal{C}$ implies $(\sup_{\mathcal{C}, \sigma}(\{t\}), \sup_{\mathcal{C}, \sigma}(\{s\})) \in \mathcal{C}(\mathcal{R})$ by definition of $\mathcal{C}(\mathcal{R})$. It is enough to prove that $\mathcal{C}(\mathcal{R}) \subseteq \mathcal{R}$.

Regarding the case $\mathcal{R} = \leq$, let $(t, s) \in \mathcal{C}(\leq)$. We reason by cases on the structure of t and s , taking into account Lemma C.4(1).

- $t = [t']^{\forall_{i \in I} \kappa_i \sigma, \forall_{i \in I} \kappa'_i \sigma}$ and $s = [s']^{\forall_{j \in J} \kappa''_j \sigma, \forall_{j \in J} \kappa'''_j \sigma}$ where $\text{cls}_{\leq, \mathcal{C}}(T) = \{[t_i]^{\kappa_i, \kappa'_i}\}_{i \in I}$ and $\text{cls}_{\leq, \mathcal{C}}(S) = \{[s_j]^{\kappa''_j, \kappa'''_j}\}_{j \in J}$ and $\emptyset \neq I \subseteq J$. We deduce $\forall_{i \in I} \kappa_i \sigma \leq \forall_{j \in J} \kappa''_j \sigma$ and $\forall_{i \in I} \kappa'_i \sigma \leq \forall_{j \in J} \kappa'''_j \sigma$. We have $t_i =_c s_j \in \mathcal{C}$ for every $i \in I$ and $j \in J$, therefore $t' = s'$ by Lemma C.4(2). We conclude $t \leq s$.
- $t = [t']^{0,0}$ and $s = [s']^{\forall_{j \in J} \kappa''_j \sigma, \forall_{j \in J} \kappa'''_j \sigma}$ where $\text{cls}_{\sim, \mathcal{C}}(T) = \{[t_i]^{\kappa_i, \kappa'_i}\}_{i \in I}$ and $\text{cls}_{\leq, \mathcal{C}}(S) = \{[s_j]^{\kappa''_j, \kappa'''_j}\}_{j \in J}$ and $\emptyset \neq I, J$. We conclude $t \leq s$ using Lemma C.4(2) like in the previous case.
- $t = [t']^{0,0}$ and $s = [s']^{0,0}$ where $\text{cls}_{\sim, \mathcal{C}}(T) = \{[t_i]^{\kappa_i, \kappa'_i}\}_{i \in I}$ and $\text{cls}_{\sim, \mathcal{C}}(S) = \{[s_j]^{\kappa''_j, \kappa'''_j}\}_{j \in J}$ and $\emptyset \neq I, J$. We conclude $t = s$ using Lemma C.4(2) like in the previous cases.
- $t = t_1 \odot t_2$ and $s = s_1 \odot s_2$ where $t_1 = \sup_{\mathcal{C}, \sigma}(\{t_i\}_{i \in I})$ and $t_2 = \sup_{\mathcal{C}, \sigma}(\{t'_i\}_{i \in I})$ and $s_1 = \sup_{\mathcal{C}, \sigma}(\{s_j\}_{j \in J})$ and $s_2 = \sup_{\mathcal{C}, \sigma}(\{s'_j\}_{j \in J})$ and $\text{cls}_{\leq, \mathcal{C}}(T) = \{t_i \odot t'_i\}_{i \in I}$ and $\text{cls}_{\leq, \mathcal{C}}(S) = \{s_j \odot s'_j\}_{j \in J}$ and $\emptyset \neq I \subseteq J$. We conclude $(t_1, s_1) \in \mathcal{C}(\leq)$ and $(t_2, s_2) \in \mathcal{C}(\leq)$ by definition of $\mathcal{C}(\leq)$.
- $t = t_1 \odot t_2$ and $s = s_1 \odot s_2$ where $t_1 = \text{zero}_{\mathcal{C}, \sigma}(\{t_i\}_{i \in I})$ and $t_2 = \text{zero}_{\mathcal{C}, \sigma}(\{t'_i\}_{i \in I})$ and $s_1 = \sup_{\mathcal{C}, \sigma}(\{s_j\}_{j \in J})$ and $s_2 = \sup_{\mathcal{C}, \sigma}(\{s'_j\}_{j \in J})$ and $\text{cls}_{\sim, \mathcal{C}}(T) = \{t_i \odot t'_i\}_{i \in I}$ and $\text{cls}_{\leq, \mathcal{C}}(S) = \{s_j \odot s'_j\}_{j \in J}$ and $\emptyset \neq I, J$. We conclude $(t_1, s_1) \in \mathcal{C}(\leq)$ and $(t_2, s_2) \in \mathcal{C}(\leq)$ by definition of $\mathcal{C}(\leq)$.
- $t = t_1 \odot t_2$ and $s = s_1 \odot s_2$ where $t_1 = \text{zero}_{\mathcal{C}, \sigma}(\{t_i\}_{i \in I})$ and $t_2 = \text{zero}_{\mathcal{C}, \sigma}(\{t'_i\}_{i \in I})$ and $s_1 = \text{zero}_{\mathcal{C}, \sigma}(\{s_j\}_{j \in J})$ and $s_2 = \text{zero}_{\mathcal{C}, \sigma}(\{s'_j\}_{j \in J})$ and $\text{cls}_{\sim, \mathcal{C}}(T) = \{t_i \odot t'_i\}_{i \in I}$ and $\text{cls}_{\sim, \mathcal{C}}(S) = \{s_j \odot s'_j\}_{j \in J}$ and $\emptyset \neq I, J$. We conclude $(t_1, s_1) \in \mathcal{C}(\leq)$ and $(t_2, s_2) \in \mathcal{C}(\leq)$ by definition of $\mathcal{C}(\leq)$.
- $t = s = \text{int}$. We conclude $t \leq s$.

The case $\mathcal{R} = \asymp$ is structurally similar. We omit the details of the proof, just observing that 0 is compatible with every use and that compatibility is closed with respect to \leq . \square

Theorem C.1 (Theorem 4.1). *If the algorithm returns σ , then σ is a minimal solution for \mathcal{C} .*

Proof. The fact that σ is a solution for \mathcal{C} follows from Lemma C.5. The solution is also minimal because so is the solution of the use constraints in \mathcal{C} picked by the algorithm at step 2 (see Table 4) and $\sup_{\mathcal{C}, \sigma}$ computes least upper bounds of uses. \square