

# Type Reconstruction Algorithms for Deadlock-Free and Lock-Free Linear $\pi$ -Calculi

Luca Padovani, Tzu-Chun Chen, Andrea Tosatto

## ▶ To cite this version:

Luca Padovani, Tzu-Chun Chen, Andrea Tosatto. Type Reconstruction Algorithms for Deadlock-Free and Lock-Free Linear  $\pi$ -Calculi. 17th International Conference on Coordination Languages and Models (COORDINATION), Jun 2015, Grenoble, France. pp.83-98, 10.1007/978-3-319-19282-6\_6. hal-01774941

# HAL Id: hal-01774941 https://inria.hal.science/hal-01774941

Submitted on 24 Apr 2018  $\,$ 

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers. L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

## Type Reconstruction Algorithms for Deadlock-Free and Lock-Free Linear $\pi$ -Calculi

Luca Padovani<sup>1</sup>, Tzu-Chun Chen<sup>1,2</sup>, and Andrea Tosatto<sup>1</sup>

<sup>1</sup> Università di Torino, Italy

<sup>2</sup> Technische Universität Darmstadt, Germany and Purdue University, USA

**Abstract.** We define complete type reconstruction algorithms for two type systems ensuring deadlock and lock freedom of linear  $\pi$ -calculus processes. Our work automates the verification of deadlock/lock freedom for a non-trivial class of processes that includes interleaved binary sessions and, to great extent, multiparty sessions as well. A Haskell implementation of the algorithms is available.

#### 1 Introduction

Type systems help finding potential errors during the early phases of software development. In the context of communicating processes, typical errors are: making invalid assumptions about the nature of a received message; using a communication channel beyond its nominal capabilities. Some type systems are able to warn against subtler errors, and sometimes can even guarantee liveness properties as well. For instance, the type systems presented in [18] for the linear  $\pi$ -calculus [14] ensure well-typed processes to be deadlock and lock free. Such stronger guarantees come at the cost of a richer type structure, hence of a greater programming effort, when programmers are supposed to explicitly annotate programs with types. In this respect, *type reconstruction* becomes a most wanted tool in the programmer's toolkit: type reconstruction is the procedure that automatically synthesizes, whenever possible, the types of the entities used by a program; in particular, the types of the channels used by a communicating process. In the present work, we describe type reconstruction algorithms for the type systems presented in [18], thereby automating the static deadlock and lock freedom analysis for a non-trivial class of communicating processes.

A *deadlock* is a configuration with pending communications that cannot complete. A paradigmatic example of deadlock modeled in the  $\pi$ -calculus is illustrated below

$$(va,b)(a?(x).b!x | b?(y).a!y)$$
 (1.1)

where the input on *a* blocks the output on *b*, and the input on *b* blocks the output on *a*. The key idea used in [18] for detecting deadlocks, which is related to earlier works by Kobayashi [11,13], is to associate each channel with a number – called *level* – specifying the relative order in which different channels should be used. In (1.1), this mechanism requires *a* to have smaller level than *b* in the left subprocess, and greater level than *b* in the right one. Since no level assignment can simultaneously satisfy both requirements, (1.1) is flagged as ill typed. This mechanism does not prevent *locks*, namely

configurations where some communication remains pending although the process as a whole can make progress. A deadlock-free configuration that is not lock free is

$$(va)(*c?(x).c!x | c!a | a!42)$$
 (1.2)

where the communication pending on *a* cannot complete. There are no interleaved communications on different channels in (1.2), therefore the level-based mechanism spots no apparent issue. The idea put forward in [18] to reject (1.2) is to also associate each channel with another number – called *ticket* – specifying the maximum number of times the channel can travel in a message. With this mechanism in place, (1.2) is ill typed because *a* would need an infinite number of tickets to travel infinitely many times on *c*.

Finding appropriate level and tickets for the channels used by a process can be difficult. We remedy to such difficulty with three contributions. First, we develop complete type reconstruction algorithms for the type systems in [18] so that appropriate level and tickets are synthesized automatically, whenever possible. The linear  $\pi$ -calculus [14], for which the type systems are defined, can model a variety of communicating systems with both static and dynamic network topologies. In particular, binary sessions [5] and, to a large extent, also multiparty sessions [18, technical report], can be encoded in it. Second, we purposely use a variant of the linear  $\pi$ -calculus with pairs instead of a polyadic calculus. While this choice has a cost in terms of technical machinery, it allows us to discuss how to deal with structured data types, which are of primary importance in concrete languages but whose integration in linear type systems requires some care [19]. We give evidence that our algorithms scale easily to other data types, including disjoint sums and polymorphic variants. Third, we present the algorithms assuming the existence of type reconstruction for the linear  $\pi$ -calculus [9,19]. This approach has two positive upshots: (1) we focus on the aspects of the algorithms concerning deadlock and lock freedom, thereby simplifying their presentation and the formal study of their properties; (2) we show how to combine in a modular way increasingly refined type reconstruction stages and how to address some of the issues that may arise in doing so.

In what follows we review the linear  $\pi$ -calculus with pairs (Section 2) and the type systems for deadlock and lock freedom of [18] (Section 3). Such type systems are unsuitable to be used as the basis for type reconstruction algorithms. So, we reformulate them to obtain reconstruction algorithms that are both correct and complete (Section 4). Then, we sketch an algorithm for solving the constraints generated by the reconstruction algorithms (Section 5) We conclude presenting a few benchmarks, further connections with related work, and directions of future research (Section 6).

The algorithms have been implemented and integrated in a tool for the static analysis of  $\pi$ -calculus processes. The archive with the source code of the tool, available at the page http://di.unito.it/hypha, includes a wide range of examples, of which we can discuss only one in the paper because of space constraints.

#### 2 The simply-typed linear $\pi$ -calculus with pairs

The process language we work with is the asynchronous  $\pi$ -calculus extended in two ways: (1) we generalize names to *expressions* to account for pairs and other data types; (2) we assume that names are explicitly annotated with *simple types* possibly inferred in

a previous reconstruction phase ("simple" means without level/ticket decorations). We annotate free names instead of bound names because, in a behavioral type system, each occurrence of a name may be used according to a different type. Typically, two distinct occurrences of the same linear channel are used for complementary I/O actions. We use  $m, n, \ldots$  to range over integer numbers; we use sets of *variables x*, *y*, ... and *channels a*, *b*, ...; names *u*, *v*, ... are either channels or variables; we let *polarities p*, *q*, ... range over subsets of  $\{?, !\}$ ; we abbreviate  $\{?\}$  with ?,  $\{!\}$  with !, and  $\{?, !\}$  with #. *Processes P*, *Q*, ..., *expressions* e, f, ..., and *simple types t*, *s*, ... are defined below:

```
ProcessP,Q::=\mathbf{0}| e?(x) \cdot P || e!f | P | Q || (va)P |*PExpressione,f::=n | u^t |(e,f) || fst(e) |snd(e)Simple typet,s::=int | p[t] | p[t]^* | t \times s
```

Expressions include integer constants, names, pairs, and the two pair projection operators fst and snd. Simple types are the regular, possibly infinite terms built using the rightmost productions in grammar above and include the type int of integers, the type p[t] of linear channels to be used according to the polarity p and carrying messages of type t, the type  $p[t]^*$  of unlimited channels to be used according to the polarity pand carrying messages of type t, and the type  $t \times s$  of pairs whose components have respectively type t and s. Recall that linear channels are meant to be used for *one* communication, whereas unlimited channels can be used any number of communications. We require every infinite branch of a type to contain infinitely many occurrences of channel constructors. For example, the term t satisfying the equation t = ?[t] is a valid type while the one satisfying the equation  $t = t \times int$  is not. We impose this requirement to simplify the formal development, but it can be lifted (for example, the implementation supports ordinary recursive types such as lists and trees).

Since we are only concerned with type reconstruction, we do not give an operational semantics of the calculus. The interested reader may refer to [14,19] for generic properties of the linear  $\pi$ -calculus and to [18] for the formalization of (dead)lock freedom. We conclude this section with a comprehensive example that is representative of a class of processes for which our type systems are able to prove deadlock and lock freedom.

Example 2.1 (full duplex communication). The term

```
*c?(x).(va)(fst(x)!a | snd(x)?(y).c!(a,y)) | c!(e,f) | c!(f,e)
```

(where we have omitted simple type annotations) models a system composed of two neighbor processes connected by channels e and f. The process spawned by c!(e, f) uses e for sending a message to the neighbor. Simultaneously, it waits on f for a message from the neighbor. The process spawned by c!(f, e) does the opposite. Each exchanged message consists of a payload (omitted) and a *continuation channel* on which subsequent messages are exchanged. Above, each process sends and receives a fresh continuation a. Once the two communications have been performed, each process iterates with a new pair of corresponding continuations.

#### **3** Type systems for deadlock and lock freedom

In this section we review the type systems ensuring deadlock and lock freedom [18] for which we want to define corresponding reconstruction algorithms. Both type systems rely on refined linear channel types of the form  $p[t]_m^n$  where the decorations n and m are respectively the *level* and the *tickets* of a channel with this type. Intuitively, levels are used for imposing an ordering on the input/output operations performed on channels: channels with lower level must be used *before* channels with higher level; tickets limit the number of "travels" for channels: a channel with m tickets can be sent at most m times in a message. From now on, we use  $T, S, \ldots$  to range over *types*, which have the same structure and constructors as simple types, but where linear channel types are decorated with levels and tickets. We write  $\lfloor T \rfloor$  for the *stripping* of T, namely for the simple type obtained by removing all level and ticket decorations from T. For example,  $\lfloor ?[int \times ![int]_m^n]^* \rfloor = ?[int \times ![int]]^*$ . Note that  $\lfloor \cdot \rfloor$  is a non-injective function.

We need some auxiliary operators. First, we extend the notion of level from channel types to arbitrary types. The level of a type T, written |T|, is an element of the set  $\mathbb{Z} \cup \{\perp, \top\}$  ordered in the obvious way and formally defined thus:

$$|T| \stackrel{\text{def}}{=} \begin{cases} \perp & \text{if } T = p[S]^* \text{ and } ? \in p \\ n & \text{if } T = p[S]_m^n \text{ and } p \neq \emptyset \\ \min\{|T_1|, |T_2|\} & \text{if } T = T_1 \times T_2 \\ \top & \text{otherwise} \end{cases}$$
(3.1)

As an example, we have  $|int \times ?[![int]_0^1]_0^0| = \min\{|int|, |?[![int]_0^1]_0^0|\} = \min\{\top, 0\} = 0$ . Intuitively, the level of *T* measures the inverse urgency for using values of type *T* in order to ensure (dead)lock freedom: the lowest level (and highest urgency)  $\perp$  is given to unlimited channels with input polarity, for which we want to guarantee input receptiveness; finite levels are reserved for linear channels; the highest level (and lowest urgency)  $\top$  is given to values such as numbers or channels with empty polarity whose use is not critical as far as (dead)lock is concerned. Note that |T| is well defined because every infinite branch of *T* has infinitely many channel constructors.

We also need an operator to *shift* the topmost levels and tickets in types. We define

$$\$_{m}^{n} T \stackrel{\text{\tiny def}}{=} \begin{cases} p [S]_{m+k}^{n+h} & \text{if } T = p [S]_{k}^{h} \\ (\$_{m}^{n} T_{1}) \times (\$_{m}^{n} T_{2}) & \text{if } T = T_{1} \times T_{2} \\ T & \text{otherwise} \end{cases}$$
(3.2)

so that, for example, we have  $\binom{2}{1}$  (int  $\times$ ? [! [int]]]  $\binom{0}{0}$  = int  $\times$ ? [! [int]]]  $\binom{1}{2}$ .

Next, we define an operator for *combining* the types of different occurrences of the same object. If an object is used according to type T in one part of a process and according to type S in another part, then it is used according to the type T + S overall, where T + S is inductively defined thus:

$$T+S \stackrel{\text{def}}{=} \begin{cases} \text{int} & \text{if } T = S = \text{int} \\ (T_1+S_1) \times (T_2+S_2) & \text{if } T = T_1 \times T_2 \text{ and } S = S_1 \times S_2 \\ (p \cup q) [T]_{h+k}^n & \text{if } T = p [T]_h^n \text{ and } S = q [T]_k^n \text{ and } p \cap q = \emptyset \\ (p \cup q) [T]^* & \text{if } T = p [T]^* \text{ and } S = q [T]^* \\ \text{undefined} & \text{otherwise} \end{cases}$$

**Table 1.** Typing rules for the deadlock-free (k = 0) and lock-free (k = 1) linear  $\pi$ -calculus.

Typing rules for e	xpressions		$\Gamma \vdash \mathbf{e}: t$				
	$\frac{[\text{T-INT}]}{\Gamma \vdash n: \texttt{int}} un(\Gamma)$	$\frac{[\text{t-name}]}{\Gamma, u: T \vdash u^{\lfloor T \rfloor}: T}$	$un(\Gamma)$				
$\frac{[\text{T-PAIR}]}{\Gamma_i \vdash \mathbf{e}_i : T}$ $\frac{\Gamma_i \vdash \mathbf{e}_i : T}{\Gamma_1 + \Gamma_2 \vdash (\mathbf{e}_1, \mathbf{e}_1)}$	$\frac{T_{i}^{(i=1,2)}}{r_{2}(i=1,2)} \qquad \frac{\Gamma \vdash e}{\Gamma \vdash fs}$	$\frac{:T \times S}{t(e):T} un(S)$	$\frac{[T-SND]}{\Gamma \vdash e: T \times S} un(T)$				
<b>Typing rules for processes</b> $\Gamma \vdash_k P$							
$\frac{\Gamma_{1} \vdash \mathbf{e}: \mathbf{P}[T]_{m}^{n}}{\Gamma_{1} \vdash \Gamma_{2} \vdash_{k}}$	$\frac{\Gamma_2, x: \$_0^n T \vdash_k P}{e?(x).P} n <  \Gamma_2 $	$\Big  \frac{\Gamma_{1} - \text{out}}{\Gamma_{1} \vdash \mathbf{e} : ! [T]} \frac{\Gamma_{1} \vdash \mathbf{e} : ! [T]}{\Gamma_{1} + \Gamma_{1}} \Big $	$\frac{n}{m} \frac{\Gamma_2 \vdash f: \$_k^n T}{\Gamma_2 \vdash_k e!f} \ n <  \Gamma_2 $				
$\frac{[{}^{\text{[T-IN*]}}]}{\Gamma_1 \vdash e : ?[T]^*}$	$\frac{\Gamma_2, x: T \vdash_k P}{*e?(x).P} \overline{un}(\Gamma_2)$	$\frac{[\text{T-OUT}^*]}{\Gamma_1 \vdash e: ! [T]^*}{\Gamma_1 + \Gamma_2}$	$\frac{\Gamma_2 \vdash f: \$_k^n T}{\vdash_k e!f} \perp <  \Gamma_2 $				
$\frac{[\text{T-IDLE}]}{\Gamma \vdash_k 0} un(\Gamma)$	$\frac{[T-PAR]}{\Gamma_1 \vdash_k P}  \frac{\Gamma_2 \vdash_k Q}{\Gamma_1 + \Gamma_2 \vdash_k P \mid Q}$	$\frac{[\text{T-NEW}]}{\Gamma \vdash_k (va)P}$	$\frac{P}{\Gamma} \qquad \frac{\prod_{i=1}^{[\text{T-NEW}^*]}}{\prod_{i=1}^{k} \prod_{i=1}^{k} \prod_{i=1}^{k} \frac{P}{\prod_{i=1}^{k} \prod_{i=1}^{k} \prod_{i=1}^{k} \prod_{i=1}^{k} \frac{P}{\prod_{i=1}^{k} \prod_{i=1}^{k} \prod$				

Type combination is partial and is only defined when the combined types have the same structure. In particular, channel types can be combined only if they have equal message types; linear channel types can be combined only if they have disjoint polarities and equal level. Also, the combination of two channel types has the union of their polarities and, in the case of linear channels, the sum of their tickets. For example, a channel that is used both with type ?[int]<sup>0</sup><sub>1</sub> and with type ![int]<sup>0</sup><sub>2</sub> is used overall according to the type ?[int]<sup>0</sup><sub>1</sub> = #[int]<sup>0</sup><sub>3</sub>.

Lastly, we define *type environments*  $\Gamma$ , ... as finite maps from names to types written  $u_1 : T_1, ..., u_n : T_n$ . As usual, dom( $\Gamma$ ) is the domain of  $\Gamma$  and  $\Gamma_1, \Gamma_2$  is the union of  $\Gamma_1$  and  $\Gamma_2$  when dom( $\Gamma_1$ )  $\cap$  dom( $\Gamma_2$ ) =  $\emptyset$ . We extend type combination to type environments:

$$\Gamma_1 + \Gamma_2 \stackrel{\text{def}}{=} \Gamma_1, \Gamma_2 \qquad \text{if } \operatorname{dom}(\Gamma_1) \cap \operatorname{dom}(\Gamma_2) = \emptyset$$
  
$$(\Gamma_1, u: T) + (\Gamma_2, u: S) \stackrel{\text{def}}{=} (\Gamma_1 + \Gamma_2), u: T + S$$

We let  $|\Gamma| \stackrel{\text{def}}{=} \min\{|\Gamma(u)| \mid u \in \operatorname{dom}(\Gamma)\}\$  be the level of a type environment, we write  $\operatorname{un}(\Gamma)$  if  $|\Gamma| = \top$  and  $\overline{\operatorname{un}}(\Gamma)$  if  $\operatorname{un}(\Gamma)$  and  $\Gamma$  has no top-level linear channel types. Note that  $\overline{\operatorname{un}}(\Gamma)$  is strictly stronger than  $\operatorname{un}(\Gamma)$ . For example, if  $\Gamma \stackrel{\text{def}}{=} x : \operatorname{int} \times \emptyset[\operatorname{int}]_0^0$  we have  $\operatorname{un}(\Gamma)$  but not  $\overline{\operatorname{un}}(\Gamma)$  because  $\Gamma(x)$  has a top-level linear channel type.

The type systems for deadlock and lock freedom are defined by the rules in Table 1 deriving judgments  $\Gamma \vdash e : T$  for expressions and  $\Gamma \vdash_k P$  for processes. The type system for deadlock freedom is obtained by taking k = 0, whereas the type system for lock

freedom is obtained by taking k = 1 and restricting all levels in linear channel types to be non negative. We illustrate the typing rules as we work through the typing derivation of the replicated process in Example 2.1. The interested reader may refer to the implementation or [18] for more examples and detailed descriptions of the rules.

Let *T* and *S* be the types defined by the equations  $T = [S]_0^0 \times ?[S]_0^0$  and  $S = ?[S]_1^1$ . We build the derivation bottom up, from the judgment stating that the whole process is well typed. Since the process is a replicated input, we apply [T-IN\*] thus:

$$\frac{c:?[T]^* \vdash c:?[T]^* \quad c:![T]^*, x: T \vdash_1 (va) (fst(x)!a \mid snd(x)?(y).c!(a,y))}{c: \#[T]^* \vdash_1 *c?(x).(va) (fst(x)!a \mid snd(x)?(y).c!(a,y))}$$

In applying this rule we have  $\Gamma_2 = c : ! [T]^*$  so the side condition  $\overline{un}(\Gamma_2)$  of  $[_{T-IN^*]}$  is satisfied: since a replicated input process is permanently available, its body cannot contain any free linear channel except those possibly received through the unlimited channel. The side condition  $\overline{un}(\Gamma_2)$ , which is stronger than simply  $un(\Gamma_2)$ , makes sure that a replicated input process does not contain linear channels and therefore is level polymorphic. We will see a use of this feature at the very end of the derivation. The continuation of the process gains visibility of the message *x* with type *T* and is a restriction of a linear channel *a*. Hence, the next step is an obvious application of  $[_{T-NEW}]$ :

$$\frac{c: ! [T]^*, x: T, a: \#[S]_3^1 \vdash_1 fst(x) ! a | snd(x)?(y).c!(a, y)}{c: ! [T]^*, x: T \vdash_1 (va) (fst(x) ! a | snd(x)?(y).c!(a, y))}$$
(3.3)

We guess level 1 and 3 tickets for *a*. The rationale is that *a* is a continuation channel that will be used *after* the channels in *x*, which have level 0, so *a* must have strictly positive level. Also, in Example 2.1 the channel *a* travels three times. At this point the typing derivation forks, for we deal with the parallel composition of two processes. This means that we have to split the type environment in two parts, each describing the resources used by the corresponding subprocess in (3.3). We have  $\Gamma = \Gamma_1 + \Gamma_2$  where

$$\Gamma \stackrel{\text{def}}{=} c : ! [T]^*, x : T, a : \#[S]_3^1 \qquad \qquad \Gamma_1 \stackrel{\text{def}}{=} x : ! [S]_0^0 \times \emptyset[S]_0^0, a : ? [S]_2^1 \\ \Gamma_2 \stackrel{\text{def}}{=} c : ! [T]^*, x : \emptyset[S]_0^0 \times ? [S]_0^0, a : ! [S]_1^1$$

Observe that  $\Gamma$  is split in such a way that: *c* only occurs in  $\Gamma_2$ , because it is only used in the right subprocess in (3.3); in each subprocess, the unused linear channel in the pair *x* is given empty polarity; the type of the continuation *a* has input polarity (and 2 tickets) in  $\Gamma_1$  and output polarity (and 1 ticket) in  $\Gamma_2$ . The type of *a* in  $\Gamma_1$  is the same as  $\$_1^0 S$ , and we use the latter form from now on. We complete the typing derivation for the left subprocess in (3.3) using  $\Gamma_1$  and applying [T-OUT]:

$$\frac{x:![S]_0^0 \times \emptyset[S]_0^0 \vdash \texttt{fst}(x):![S]_0^0}{x:![S]_0^0 \times \emptyset[S]_0^0, a: \$_1^0 S \vdash_1 \texttt{fst}(x)!a} 0 < |\$_1^0 S| = 1$$

The side condition 0 < 1 ensures that the message has higher level than the channel on which it travels, according to the intuition that the message can only be used *after* the communication has occurred and the message has been received. In this case, the level of fst(x) is 0 that is smaller than the level of a, which is 1. Shifting the tickets from the type of *a* consumes one of its tickets, meaning that after this communication *a* gets closer to the point where it must be the subject of a communication.

Concerning the right subprocess in (3.3), we use  $\Gamma_2$  above and apply [T-IN] to obtain

$$\frac{x: \emptyset[S]_0^0 \times ?[S]_0^0 \vdash \operatorname{snd}(x): ?[S]_0^0 \quad c: ![T]^*, a: ![S]_1^1, y: S \vdash_1 c! (a, y)}{c: ![T]^*, x: \emptyset[S]_0^0 \times ?[S]_0^0, a: ![S]_1^1 \vdash_1 \operatorname{snd}(x)?(y).c! (a, y)} 0 < 1$$

The side condition 0 < 1 checks that the level of the linear channel used for input is smaller than the level of any other channel occurring free in the continuation of the process. In this case, *c* has level  $\top$  because it is an unlimited channel with output polarity, whereas *a* has level 1. To close the derivation we must type the recursive invocation of *c*. We do so with an application of [T-OUT\*]:

$$\frac{c:![T]^* \vdash c:![T]^*}{c:![T]^*, a:![S]_1^1, y:S \vdash (a,y):\$_1^1 T} \perp < 1$$

The side condition  $\perp < 1$  ensures that no unlimited channel with input polarity is sent in the message. This is necessary to guarantee input receptiveness on unlimited channels. There is a mismatch between the actual type  $\$_1^1 T$  and the expected type T of the message. The shifting on the tickets is due, once again, to the fact that 1 ticket is required and consumed for the channels to travel. The shifting on the levels realizes a form of *level polymorphism* whereby we are allowed to send on *c* a pair of channels with level 1 even if *c* expects a pair of channels with level 0. This is safe because we know, from the side condition of [T-IN\*], that the receiver of the message does not own any linear channel except those possibly contained in the message itself. Therefore, the exact level of the channels in the message is irrelevant, as long as it is obtained by shifting of the expected message type. Level polymorphism is a key distinguishing feature of our type systems that makes it possible to deal with non-trivial recursive processes.

#### 4 Type reconstruction

We now face the problem of defining a type reconstruction algorithm for the type system presented in the previous section. The input of the algorithm is a process P where names are explicitly annotated with simple types, possibly resulting from a previous reconstruction stage [9,19]. Notwithstanding such explicit annotations, the typing rules in Table 1 rely on guesses concerning (i) the splitting of type environments, (ii) levels and tickets that decorate linear channel types, and (iii) how tickets are distributed in combined types. We address these issues using standard strategies. Concerning (i), we synthesize type environments for expressions and processes by looking at the free names occurring in them. Concerning (ii) and (iii), we proceed in two steps: first, we transform each simple type t in P into a *type expression* T that has the same structure as t, but where we use fresh level and ticket *variables* in every slot where a level or a ticket is expected; we call this transformation *dressing*. Then, we accumulate (rather than check) the constraints that these level and ticket variables should satisfy, as by the side conditions of the typing rules (Table 1). Finally, we look for a solution of these

constraints. It turns out that the accumulated constraints can always be expressed as an integer programming problem for which there exist dedicated solvers.

There is still a subtle source of ambiguity in the procedure outlined so far. We have remarked that stripping is a non-injective function, meaning that different types may be stripped to the same simple type. For example, if we take  $T = ?[T]_1^1$  and  $S = ?[T]_0^0$  we have |T| = |S| = s where s = ?[s]. Now, if we were to reconstruct either T or S from s, we would have to dress s with level and ticket variables in every slot where a level or a ticket is expected. But since s is infinite, such dressing is not unique. For example,  $T = ?[T]_{\theta_1}^{\eta_1}$  and  $S = ?[T]_{\theta_2}^{\eta_2}$  are just two of the infinitely many possible dressings of *s* with level and ticket variables: in T we have used two distinct variables  $\eta_1$  and  $\theta_1$ , one for each slot; in S we have used four. The problem is that from the dressing T we can only reconstruct T, by taking  $\eta_1 = \theta_1 = 1$ , whereas from the dressing S we can reconstruct both T (by assigning all variables to 1) as well as S, by taking  $\eta_1 = \theta_1 = 1$ and  $\eta_2 = \theta_2 = 0$ . This means that the choice of the number of integer variables we use in dressing (infinite) simple types constrains the types that we can reconstruct from them, which is a risk for the completeness of the type reconstruction algorithms. To cope with this issue, we dress simple types *lazily*, only to their topmost linear channel constructors, and we put fresh type variables in place of message types, leaving them undressed. It is only when the message is used that we (lazily) dress its type as well. The introduction of fresh type variables for message types means that we redo part of the work already carried out for reconstructing simple types [19]. This appears to be an inevitable price to pay to have completeness of the type reconstruction algorithms, when they build on top of (instead of being performed together with) previous stages.

To formalize the algorithms, we introduce countable sets of *type variables*  $\alpha$ ,  $\beta$  and of *integer variables*  $\eta$ ,  $\theta$ ; *type expressions* and *integer expressions* are defined below:

Type expression	T,S	::=	int $  \alpha   p[T]^{\lambda}_{\tau}   p[T]^*   T \times S$
Integer expression	$\lambda, arepsilon,  au$	::=	$n \mid \eta \mid \varepsilon + \varepsilon \mid \varepsilon - \varepsilon$

Type expressions differ from types in three ways: they are always finite, they have integer expressions in place of levels and tickets, and they include type variables  $\alpha$  denoting unknown types awaiting to be lazily dressed. Integer expressions are linear polynomials of integer variables.

We say that T is *proper*, written prop(T), if all the type variables in T are guarded by a channel constructor. For example, both int and  $p[\alpha]^*$  are proper (all type variables occur within channel types), but  $\alpha$  and  $int \times \alpha$  are not. Since the level and tickets of a type expression are solely determined by its top-level linear channel constructors, properness characterizes those type expressions that are "sufficiently dressed" so that it is possible to extract their level and to combine them with other type expressions, even if these type expressions contain type variables.

We now revisit and adapt all the auxiliary operators and notions defined for types to type expressions. Recall that the level of *T* is the minimum level of any topmost linear channel type in *T*, or  $\perp$  if *T* has a topmost unlimited channel type with input polarity. Since a type expression T may contain unevaluated level expressions, we cannot compute a minimum level in general. However, a quick inspection of Table 1 reveals that minima of levels always occur on the right hand side of inequalities, and an inequality like  $n < \min\{m_i \mid i \in I\}$  can equivalently be expressed as a set of inequalities

 $\{n < m_i \mid i \in I\}$ . Following this observation, we define the *level*  $|\mathsf{T}|$  of a proper type expression T as the *set* of level expressions that decorate the topmost linear channel types in T, and possibly the element  $\perp$ . Formally:

$$|\mathsf{T}| \stackrel{\text{def}}{=} \begin{cases} \{\bot\} & \text{if } \mathsf{T} = p \, [\mathsf{S}]^* \text{ and } ? \in p \\ \{\lambda\} & \text{if } \mathsf{T} = p \, [\mathsf{S}]^{\lambda}_{\tau} \text{ and } p \neq \emptyset \\ |\mathsf{T}_1| \cup |\mathsf{T}_2| & \text{if } \mathsf{T} = \mathsf{T}_1 \times \mathsf{T}_2 \\ \emptyset & \text{otherwise} \end{cases}$$
(4.1)

We write un(T) if  $|T| = \emptyset$ , in which case T denotes an unlimited type. *Shifting* for proper type expressions is defined just like for types, except that we symbolically record the sum of level/ticket expressions instead of computing it:

$$\$^{\lambda}_{\tau} \mathsf{T} \stackrel{\text{def}}{=} \begin{cases} p[\mathsf{S}]^{\lambda+\lambda'}_{\tau+\tau'} & \text{if } \mathsf{T} = p[\mathsf{S}]^{\lambda'}_{\tau'} \\ (\$^{\lambda}_{\tau} \mathsf{T}_1) \times (\$^{\lambda}_{\tau} \mathsf{T}_2) & \text{if } \mathsf{T} = \mathsf{T}_1 \times \mathsf{T}_2 \\ \mathsf{T} & \text{otherwise} \end{cases}$$
(4.2)

Because type expressions may contain type and integer variables, we cannot determine *a priori* whether the combination of two type expressions is possible. For instance, the combination of  $\left[T\right]_{\tau_1}^{\lambda_1}$  and  $\left[S\right]_{\tau_2}^{\lambda_2}$  is possible only if T and S denote the same type and if  $\lambda_1$  and  $\lambda_2$  evaluate to the same level. We cannot check these conditions right away, when T, S and the level expressions contain variables. Instead, we record these conditions into a *constraint*. Constraints  $\varphi$ , ... are conjunctions of *type constraints* T = S(equality relations between type expressions) and *integer constraints*  $\varepsilon \leq \varepsilon'$  (inequality relations between integer expressions). Formally, their syntax is defined by

Constraint 
$$\varphi$$
 ::= true | T = T |  $\varepsilon \leq \varepsilon$  |  $\phi \land \phi$ 

We write  $\varepsilon < \varepsilon'$  in place of  $\varepsilon + 1 \le \varepsilon'$  and  $\varepsilon = \varepsilon'$  in place of  $\varepsilon \le \varepsilon' \land \varepsilon' \le \varepsilon$ ; if  $\mathscr{E} = \{\varepsilon_i\}_{i \in I}$  is a finite set of integer expressions, we write  $\varepsilon < \mathscr{E}$  for the constraint  $\bigwedge_{i \in I} \varepsilon < \varepsilon_i$ ; finally, we write dom( $\varphi$ ) for the set of type expressions occurring in  $\varphi$ .

The *combination* operator  $T \sqcup S$  for type expressions returns a pair R;  $\varphi$  made of the resulting type expression R and the constraint  $\varphi$  that must be satisfied for the combination to be possible. The definition of  $\sqcup$  mimics exactly that of + in Section 3, except that all non-checkable conditions accumulate in constraints:

	$ \begin{pmatrix} \text{int} \\ (n \cup a) [T']^{\lambda} \end{bmatrix} $	; true : $T' = S' \land \lambda = \lambda'$	if $T = int$ and $S = int$ if $T = n[T']^{\lambda}$ and $S = a[S']^{\lambda'}$ .
$T \sqcup S \stackrel{\text{\tiny def}}{=} C$	$(p \cup q) [T']_{\tau + \tau'}^*$ $R_1 \times R_2$	$; \mathbf{T}' = \mathbf{S}' \\ ; \boldsymbol{\varphi}_1 \wedge \boldsymbol{\varphi}_2$	if $T = p[T']_{\tau}$ and $S = q[S']_{\tau'}$ and $p \cap q = \emptyset$ if $T = p[T']^*$ and $S = q[S']^*$ if $T = T_1 \times T_2$ and $S = S_1 \times S_2$ and $T_i \sqcup S_i = R_i; \varphi_i$
	undefined		otherwise

Like type combination, also  $\sqcup$  is a partial operator:  $T \sqcup S$  is undefined if T and S are structurally incompatible (*e.g.*, if T = int and  $S = p[int]^*$ ) or if T and S are not proper. When  $T \sqcup S$  is defined, though, the resulting type expression is always proper.

We use  $\Delta$ , ... to range over *type expression environments* (or just environments, for short), namely finite maps from names to type expressions, and we inherit all the notation introduced for type environments. We let  $|\Delta| \stackrel{\text{def}}{=} \bigcup_{u \in \text{dom}(\Delta)} |\Delta(u)|$  and write  $\overline{\text{un}}(\Delta)$  if  $|\Delta| = \emptyset$  and  $\Delta$  has no top-level linear channel type in its range. By now, the extension of  $\sqcup$  to environments is easy to imagine: when defined,  $\Delta_1 \sqcup \Delta_2$  is a pair  $\Delta$ ;  $\varphi$  made of the resulting environment  $\Delta$  and of a constraint  $\varphi$  that results from the combination of the type expressions in  $\Delta_1$  and  $\Delta_2$ . More precisely:

$$\begin{array}{ll} \Delta_1 \sqcup \Delta_2 \stackrel{\text{\tiny det}}{=} \Delta_1, \Delta_2 \ ; \mbox{true} & \mbox{if } \operatorname{dom}(\Delta_1) \cap \operatorname{dom}(\Delta_2) = \emptyset \\ (\Delta_1, u : \mathsf{T}) \sqcup (\Delta_2, u : \mathsf{S}) \stackrel{\text{\tiny def}}{=} \Delta, u : \mathsf{R} \ ; \ \varphi \land \varphi' & \mbox{if } \Delta_1 \sqcup \Delta_2 = \Delta; \varphi \ \mbox{and } \mathsf{T} \sqcup \mathsf{S} = \mathsf{R}; \varphi' \end{array}$$

The last notion we need to formalize, before introducing the reconstruction algorithms, is that of *dressing*. Dressing a simple type t means placing fresh integer variables in the level/ticket slots of t. Formally, we say that T is a *dressing* of t if  $t \uparrow T$  is inductively derivable by the following rules which pick globally fresh variables:

$$\texttt{int} \uparrow \texttt{int} \qquad \frac{\alpha \text{ fresh}}{p[t]^* \uparrow p[\alpha]^*} \qquad \frac{\alpha, \eta, \theta \text{ fresh}}{p[t] \uparrow p[\alpha]_{\theta}^{\eta}} \qquad \frac{t_i \uparrow \mathsf{T}_i^{(i=1,2)}}{t_1 \times t_2 \uparrow \mathsf{T}_1 \times \mathsf{T}_2}$$

Note that the decoration of t with fresh integer variables stops at the topmost channel types in t and that message types are left undecorated. By definition, the dressing of a simple type is always a proper type expression.

We can now present the type reconstruction algorithms, defined by the rules in Table 2. The rules in the upper part of the table derive judgments of the form  $e: T \triangleright \Delta; \varphi$ , stating that e has type T in the environment  $\Delta$  if the constraint  $\varphi$  is satisfied. The expression e is the only "input" of the judgment, while T,  $\Delta$ , and  $\varphi$  are synthesized from it. There is a close correspondence between these rules and those for expressions in Table 1. Observe the use of  $\sqcup$  where + was used in Table 1, the accumulation of constraints from the premises to the conclusion of each rule and, most notably, the dressing of the simple type that annotates u in [I-NAME]. Type expressions synthesized by the rules are always proper, so the side conditions in [I-FST] and [I-SND] can be safely checked.

The rules in the lower part of the table derive judgments of the form  $P \triangleright_k \Delta; \varphi$ , stating that *P* is well typed in the environment  $\Delta$  if the constraint  $\varphi$  is satisfied. The parameter *k* plays the same role as in the type system (Table 1). The process *P* and the parameter *k* are the only "inputs" of the judgments, and  $\Delta$  and  $\varphi$  are synthesized from them. All rules except [I-WEAK] have a corresponding one in Table 1. Like for expressions, environments are combined through  $\sqcup$  and constraints accumulate from premises to conclusions. We focus on the differences with respect to the typing rules.

In rule [T-IN], the side condition verifies that the level of the channel e on which an input is performed is smaller than the level of any channel used for typing the continuation process *P*. This condition can be decomposed in two parts: (1) no unlimited channel with input polarity must be in *P*; this condition is necessary to ensure input receptiveness on unlimited channels in the original type system [18] and is expressed in [I-IN] as the side condition  $\perp \notin |\Delta_2|$ , which can be checked on type expressions environments directly; (2) the level of e must satisfy the ordering with respect to all the linear channels in *P*; this is expressed in [I-IN] as the constraint  $\lambda < |\Delta_2|$ , where  $\lambda$  is the level of e. The same side condition and constraint are found in [I-OUT].

Reconstruction rules for ex	pressions	$e:T\blacktriangleright\Delta;\varphi$			
[I-INT] $n: int \triangleright \emptyset; true$	$\frac{\mathbf{e}_{i}:T_{i} \blacktriangleright \Delta_{i}; \boldsymbol{\varphi}_{i} \stackrel{(i=1,2)}{(\mathbf{e}_{1},\mathbf{e}_{2}):T_{1} \times T_{2} \blacktriangleright \Delta; \bigwedge_{1 \leq i}}$	$_{i\leq 3} \varphi_i \ \Delta_1 \sqcup \Delta_2 = \Delta; \varphi_3$			
$\frac{[I-NAME]}{u^t: T \blacktriangleright u: T; true} t \uparrow T$	$\frac{e: T \times S \blacktriangleright \Delta; \varphi}{\mathtt{fst}(e): T \blacktriangleright \Delta; \varphi} \ un(S)$	$\frac{e: T \times S \blacktriangleright \Delta; \varphi}{\operatorname{snd}(e): S \blacktriangleright \Delta; \varphi}  un(T)$			
Reconstruction rules for pr	ocesses	$P \blacktriangleright_k \Delta; \varphi$			
$\frac{P \blacktriangleright_k \Delta; \varphi}{P \blacktriangleright_k \Delta, u: T; \varphi}  un(T)$	$\frac{[I-IDLE]}{0 \blacktriangleright_k 0; true} \qquad \frac{P_i \blacktriangleright_k \Delta}{P_1 \mid P_2 \blacktriangleright_k}$	$\frac{i;\varphi_i^{(i=1,2)}}{\Delta;\bigwedge_{1\leq i\leq 3}\varphi_i} \ \Delta_1\sqcup\Delta_2=\Delta;\varphi_3$			
$\frac{e:?[T]_{\tau}^{\lambda} \blacktriangleright \Delta_{1};\varphi_{1} \qquad P \blacktriangleright_{k} \Delta_{2}, x:S;\varphi_{2}}{e?(x).P \blacktriangleright_{k} \Delta; \bigwedge_{1 \le i \le 3} \varphi_{i} \land T = \$_{0}^{-\lambda} S \land \lambda <  \Delta_{2} } \qquad \bot \notin  \Delta_{2}  \\ \Delta_{1} \sqcup \Delta_{2} = \Delta;\varphi_{3}$					
$\frac{e: ! [T]_{\tau}^{\lambda} \blacktriangleright \Delta_{1}; \varphi_{1}}{e! f \blacktriangleright_{k} \Delta; \bigwedge_{1 \leq i \leq 3} \varphi_{i} \land T = \$}$	$ \begin{array}{l} \mathbf{f}: \mathbf{S} \blacktriangleright \Delta_2; \boldsymbol{\varphi}_2 \\ \mathbf{g}_{-k}^{-\lambda}  \mathbf{S} \land \lambda <  \Delta_2  \\ \end{array} \xrightarrow{\perp \notin  \Delta_2 } \Delta_1 \sqcup \Delta_2 = \Delta  \end{array} $	$\sum_{\lambda;\varphi_3}^{[\text{I-NEW}]} \frac{P \blacktriangleright_k \Delta, a: \#[T]_{\tau}^{\lambda}; \varphi}{(va) P \blacktriangleright_k \Delta; \varphi}$			
$\frac{\stackrel{[I-IN^*]}{=:?[T]^* \blacktriangleright \Delta_1; \varphi_1 \qquad P \blacktriangleright_k \Delta_2, x: S; \varphi_2}{*e?(x) \cdot P \blacktriangleright_k \Delta; \bigwedge_{1 \le i \le 3} \varphi_i \land T = S}  \overline{\Delta_1} \sqcup \Delta_2 = \Delta; \varphi_3$					
$\frac{e:![T]^* \blacktriangleright \Delta_1; \varphi_1}{e!f \blacktriangleright_k \Delta; \bigwedge_{1 \le i \le 3} \varphi_i \land}$	$\frac{f:S \blacktriangleright \Delta_2; \varphi_2}{T = \$_{-k}^{-\eta} S} \begin{array}{c} \bot \notin  \Delta_2  \\ \Delta_1 \sqcup \Delta_2 = \Delta; \varphi_1 \\ \eta \text{ fresh} \end{array}$	$3 \qquad \frac{P \blacktriangleright_k \Delta, a : \#[T]^*; \varphi}{(va) P \blacktriangleright_k \Delta; \varphi}$			

Table 2. Type reconstruction rules for expressions and processes.

In [T-IN], [T-OUT], and [T-OUT\*], shifting is used for updating message levels, consuming tickets, and realizing level polymorphism. In rules [I-IN], [I-OUT], and [I-OUT\*], analogous shiftings are performed on type expressions, except that they are inverted and recorded in constraints. For example, when typing the continuation *P* of a process e?(x).*P* using [T-IN], if e has type ?[T]<sup>*n*</sup><sub>*m*</sub> then the type of *x* is required to be  $\$_0^n T$ . In the reconstruction algorithm, we record this requirement as the constraint  $T = \$_0^{-\lambda} S$ , where S is the type synthesized for *x* in *P*. We invert the shifting because shifting is defined only on proper type expressions, and in [I-IN] (and the other rules mentioned) only S is guaranteed to be proper, while T in general is not.

Finally, note that [I-WEAK] has no correspondent rule in Table 1. This rule is necessary because the premises of [I-IN], [I-IN\*], [I-NEW], and [I-NEW\*] *assume* that bound names occur in their scope. Since type environments are generated by the algorithm as it works through an expression or a process, this assumption may not hold if a bound name is never used in its scope. Naturally, the type T of an unused name must be unlimited, whence the constraint un(T). We also require T to be proper, to preserve the invariant that all environments synthesized by the algorithms have proper types. In principle, [I-WEAK] makes the rule set in Table 2 not syntax directed, which is a problem if we want to consider this as an algorithm. In practice, the places where [I-WEAK] may be necessary are easy to spot (in the premises of all the aforementioned rules for the binding constructs). What we gain with [I-WEAK] is a simpler presentation of the rules.

To state the properties of the reconstruction algorithm, we need a notion of constraint satisfiability. A *variable assignment*  $\sigma$  is a map from type/integer variables to types/integers. We say that  $\sigma$  covers X if  $\sigma$  provides assignments to all the type/integer variables occurring in X, where X may be a constraint, a type/integer expression, or an environment. When  $\sigma$  covers X, the *application* of  $\sigma$  to X, written  $\sigma$ X, substitutes all type/integer variables according to  $\sigma$  and evaluates all integer expressions in X. When  $\sigma$  covers  $\varphi$ , we say that  $\sigma$  satisfies  $\varphi$  if  $\sigma \models \varphi$  is derivable by the rules:

$$\frac{1}{\sigma \vDash \mathsf{True}} \qquad \frac{1}{\sigma \vDash \mathsf{T} = \mathsf{S}} \quad \sigma \mathsf{T} = \sigma \mathsf{S} \qquad \frac{1}{\sigma \vDash \varepsilon \le \varepsilon'} \quad \sigma \varepsilon \le \sigma \varepsilon' \qquad \frac{1}{\sigma \vDash \varphi_i} \frac{1}{(i=1,2)} \frac{1}{\sigma \vDash \varphi_i} \frac{1}{\sigma$$

Whenever we apply an assignment  $\sigma$  to a set of type expressions in reference to a derivation that is parametric on *k*, we will implicitly assume that all integer expressions in ticket slots evaluate to non-negative integers and that, if k = 1, all integer expressions in level slots evaluate to non-negative integers. The value of *k* and the set of type expressions will always be clear from the context.

The reconstruction algorithm is *correct*, namely each derivation obtained through the algorithm such that the resulting constraint is satisfiable corresponds to a derivation in the type system:

#### **Theorem 4.1** (correctness). If $P \triangleright_k \Delta$ ; $\varphi$ and $\sigma \vDash \varphi$ and $\sigma$ covers $\Delta$ , then $\sigma \Delta \vdash_k P$ .

The algorithm is also *complete*, meaning that if there exists a typing derivation for the judgment  $\Gamma \vdash_k P$ , then the algorithm is capable of synthesizing an environment  $\Delta$  from which  $\Gamma$  can be obtained by means of a suitable variable assignment:

**Theorem 4.2 (completeness).** *If*  $\Gamma \vdash_k P$ , *then*  $P \triangleright_k \Delta$ ;  $\varphi$  *for some*  $\Delta$ ,  $\varphi$ , *and*  $\sigma$  *such that*  $\sigma \vDash \varphi$  *and*  $\Gamma = \sigma \Delta$ .

Note that the above results do not give any information about how to verify whether there exists a  $\sigma$  such that  $\sigma \vDash \varphi$  and, in this case, how to find such  $\sigma$ . These problems will be addressed in Section 5. We conclude this section showing the reconstruction algorithm at work on the replicated process in Example 2.1.

*Example 4.1.* Below is the replicated process in Example 2.1, where we have numbered and named the relevant rules used by the algorithm as it visits the process bottom-up, left-to-right: (4) [I-PAR]

$$\underbrace{*c?(x).(va)(\operatorname{fst}(x)!a|}_{(1)[I-OUT]} \underbrace{\operatorname{snd}(x)?(y)}_{(3)[I-IN]} \underbrace{c!(a,y)}_{(2)[I-OUT^*]}$$

Table 3. Type environment and constraints generated for the process in Example 2.1.

i	С	x	а	у	Constraint
(1)		$! [\alpha_1]_{\theta_1}^{\eta_1} \times \emptyset[\alpha_2]_{\theta_2}^{\eta_2}$	? $[\alpha_3]_{\theta_3}^{\eta_3}$		$\alpha_1 = ? [\alpha_3]_{\theta_3 - k}^{\eta_3 - \eta_1} \land \eta_1 < \eta_3$
(2)	![α <sub>4</sub> ]*		$[\alpha_{5}]_{\theta_{5}}^{\eta_{5}}$	$\left[\alpha_{6}\right]_{\theta_{6}}^{\eta_{6}}$	$\left \alpha_{4}=!\left[\alpha_{5}\right]_{\theta_{5}-k}^{\eta_{5}-\eta_{4}}\times?\left[\alpha_{6}\right]_{\theta_{6}-k}^{\eta_{6}-\eta_{4}}\right $
(3)		$\emptyset[lpha_7]^{\eta_7}_{\theta_7} imes$ ? $[lpha_8]^{\eta_8}_{\theta_8}$			$lpha_8=$ ? [ $lpha_6$ ] $rac{\eta_6-\eta_8}{ heta_6}\wedge\eta_8<\eta_5$
(4)		$! [\alpha_1]_{\theta_1+\theta_7}^{\eta_1} \times ? [\alpha_2]_{\theta_2+\theta_8}^{\eta_2}$	$\#[\alpha_3]_{\theta_3+\theta_5}^{\eta_3}$		$\alpha_1 = \alpha_7 \wedge \alpha_2 = \alpha_8 \wedge \alpha_3 = \alpha_5$
					$\wedge \eta_1 = \eta_7 \wedge \eta_2 = \eta_8 \wedge \eta_3 = \eta_5$
(5)	#[α <sub>9</sub> ]*				$\alpha_9 = ! [\alpha_1]_{\theta_1+\theta_7}^{\eta_1} \times ? [\alpha_2]_{\theta_2+\theta_8}^{\eta_2}$
					$\wedge \alpha_9 = \alpha_4$

Table 4. Constraint entailment rules.

$\frac{[\text{S-LEVEL}]}{\varphi \vdash_1 0 \leq \lambda} \ p[\texttt{T}]^{\lambda}_{\tau} \in$	$dom(\varphi) \qquad \frac{[\text{S-TICKET}]}{\varphi \vdash_k 0 \leq}$	$\frac{[\text{S-TICKET}]}{\varphi \vdash_k 0 \leq \tau} \ p[T]^{\lambda}_{\tau} \in dom(\varphi)$		
$\frac{\text{[s-conj]}}{\varphi_1 \land \varphi_2 \vdash_k \varphi_i}  i \in \{1,2\}$	$\frac{\varphi \vdash_k T = S}{\varphi \vdash_k S = T} \qquad \frac{\varphi \vdash_k T}{\varphi \vdash_k S = T}$	$\frac{[RANS]}{\varphi \vdash_k T = R} \qquad \varphi \vdash_k R = S$		
$\frac{\varphi \vdash_k p [T]_{\tau_1}^{\lambda_1} = p [S]_{\tau_2}^{\lambda_2}}{\varphi \vdash_k T = S \land \lambda_1 = \lambda_2 \land \tau_1 = \tau_2}$	$\frac{\varphi \vdash_{k} p[T]^{*} = p[S]}{\varphi \vdash_{k} T = S}$	$\frac{\varphi \vdash_k T_1 \times T_2 = S_1 \times S_2}{\varphi \vdash_k T_1 = S_1 \wedge T_2 = S_2}$		

Each subprocess triggers one rule of the reconstruction algorithm which synthesizes a type environment and possibly generates some constraints. Table 3 summarizes the parts of the environments and the constraints produced at each step of the reconstruction algorithm with parameter k. We have omitted the step concerning the restriction on a, which just removes a from the environment and introduces no constraints.

#### 5 Constraint solving

We sketch an algorithm that determines whether a constraint  $\varphi$  is satisfiable and, in this case, computes an assignment that satisfies it. The presentation is somewhat less formal since the key steps of the algorithm are instances of well-known techniques. The algorithm is structured in three phases, *saturation*, *verification*, and *synthesis*.

The constraint  $\varphi$  produced by the reconstruction algorithm does not necessarily mention all the relations that must hold between integer variables. For example, the constraint  $\eta_3 - \eta_1 = \eta_6 - \eta_8 \land \theta_3 - k = \theta_6$  is implied by those in Table 3, but it appears nowhere. Finding all the integer constraints entailed by a given  $\varphi$ , regardless of whether such constraints are implicit or explicit, is essential because we use an external solver for solving them. The aim of the *saturation phase* is to find all such integer constraints. Table 4 defines an inference system for deriving entailments  $\varphi \vdash_k \varphi'$ . The parameter *k* plays the same role as in the type system. Rules [s-level] and [s-ticket] introduce nonnegativity constraints for integer expressions that occur in level and ticket slots; level expressions are required to be non-negative only for lock freedom analysis, when k = 1; rule [s-conj] decomposes conjunctions; rules [s-symm] and [s-trans] compute the symmetric and transitive closure of type equality; finally, [s-chan], [s-chan\*], and [s-pair] state expected congruence rules. We let  $\hat{\varphi} \stackrel{\text{def}}{=} \bigwedge_{\varphi \vdash_k \varphi'} \varphi'$ . Clearly  $\hat{\varphi}$  can be computed in finite time and is satisfiable by the same assignments as (*i.e.*, it is equivalent to)  $\varphi$ .

The verification phase checks whether  $\hat{\varphi}$  is satisfiable and, in this case, computes an assignment  $\sigma_{int}$  that satisfies the integer constraints in it. In  $\hat{\varphi}$  all the integer constraints are explicit. These are typical constraints of an integer programming problem, for which it is possible to use dedicated (complete) solvers that find a  $\sigma_{int}$  when it exists (our tool supports GLPK<sup>3</sup> and lpsolve<sup>4</sup>). When this is the case, the type constraints in  $\hat{\varphi}$  are satisfiable if, for each type constraint of the form T = S, either T or S are type variables, or T and S have the same topmost constructor, *i.e.* they are either both int, or both unlimited/linear channel types with the same polarity, or both product types.

The *synthesis phase* computes an assignment that satisfies  $\varphi$ . This is found by applying  $\sigma_{int}$  to all the type constraints in  $\hat{\varphi}$ , by choosing a canonical constraint of the form  $\alpha = T$  where T is proper for each  $\alpha \in \text{dom}(\hat{\varphi})$ , and then by solving the resulting system  $\{\alpha_i = T_i\}$  of equations. By [4, Theorem 4.2.1], this system has exactly one solution  $\sigma_{type}$  and now  $\sigma_{int} \cup \sigma_{type} \models \varphi$ . There may be type variables  $\alpha$  for which there is no  $\alpha = T$  constraint with T proper. These type variables denote values not used by the process, like a message that is received from one channel and just forwarded on another one. These variables are assigned a type that can be computed canonically.

*Example 5.1.* The constraints shown in Table 3 entail  $0 \le \theta_3 - k$  and  $0 \le \theta_5 - k$  and  $0 \le \theta_6 - k$  namely  $k \le \theta_3$  and  $k \le \theta_5$  and  $k \le \theta_6$  must hold. When k = 0, these constraints can be trivially satisfied by assigning 0 to all ticket variables. When k = 1, from the type of *a* at step (4) of the reconstruction algorithm we deduce that *a* must have at least 2 tickets. Indeed, *a* is sent in two messages. It is only considering the remaining processes c! (e, f) and c! (f, e) that we learn that *y* is instantiated with *a*. Then, *a* needs one more ticket, to account for the further and last travel in the recursive invocation c! (a, y).

#### 6 Concluding remarks

A key distinguishing feature of the type systems in [18] is the use of polymorphic recursion. Type reconstruction in presence of polymorphic recursion is notoriously undecidable [10,7]. In our case, polymorphism solely concerns levels and reconstruction turns out to be doable. A similar situation is known for *effect systems* [1], where polymorphic recursion restricted to effects does not prevent complete type reconstruction [2].

We have conducted some benchmarks on generalizations of Example 2.1 to Ndimensional hypercubes of processes using full-duplex communication. The table below reports the reconstruction times for the analysis of an hypercube of side 5 and N varying from 1 to 4. The table details the dimension, the number of processes and channels, and the times (in seconds) spent for linearity analysis [19], constraint generation (Section 4) and saturation, solution of level and ticket constraints (Section 5). The

<sup>&</sup>lt;sup>3</sup> http://www.gnu.org/software/glpk/

<sup>&</sup>lt;sup>4</sup> http://sourceforge.net/projects/lpsolve/

solver used for level and ticket constraints is GLPK 4.48 and times were measured on a 13" MacBook Air running a 1.8 GHz Intel Core i5 with 4 GB of 1600MHz DDR3.

	Ν	Processes	Channels	Linearity	Gen.+Sat.	Levels	Tickets	Overall
ĺ	1	5	8	0.021	0.006	0.002	0.003	0.032
	2	25	80	0.128	0.051	0.009	0.012	0.200
İ	3	125	600	1.439	0.844	0.069	0.124	2.477
ĺ	4	625	4000	33.803	26.422	1.116	3.913	65.254

Reconstruction times scale almost linearly in the number of channels as long as there is enough free main memory. With N = 4, however, the used memory exceeds 10GB causing severe memory (de)compression and swapping. The running time inflates consequently. We have not determined yet the precise causes of such disproportionate consumption of memory, which the algorithms do not seem to imply. We suspect that they are linked to our naive implementation of the algorithms in a lazy language (Haskell), but a more rigorous profiling analysis is left for future investigation. Integer programming problems are NP-hard in general, but the time used for integer constraint resolution appears negligible compared to the other phases. As suggested by one reviewer, the particular nature of such constraints indicates that there might be more clever way of solving them, for example by using SMT solvers.

Our work has been inspired by previous type systems ensuring (dead)lock freedom for generic  $\pi$ -calculus processes [11,13] and corresponding type reconstruction algorithms [12]. These type systems and ours are incomparable: [11,13] use sophisticated behavioral types that provide better accuracy with respect to unlimited channels as used for modeling mutual exclusion and concurrent objects. On the other hand, our type systems exploit level polymorphism for dealing with recursive processes in cyclic topologies, often arising in the modeling of parallel algorithms and sessions. Whether and how the strengths of both approaches can be combined together is left for future research. A more thorough comparison between these works can be found in [18].

There is a substantial methodological difference between our approach and those addressing sessions, particularly *multiparty sessions* [8,6]. Session-based approaches are top down and type driven: types/protocols come first, and are used as a guidance for developing programs that follow them. These approaches guarantee by design a number of properties, among which (dead)lock freedom when different sessions are not interleaved. Our approach is bottom up and program driven: programs come first, and are used for inferring types/protocols. The two approaches can integrate and complement each other. For example, type reconstruction may assist in the verification of legacy or third-party code (for which no type information is available) or for checking the impact of code changes due to refactoring and/or debugging. Also, some protocols are hard to describe a priori. For example, describing the essence of full-duplex communications (Example 2.1) is far from trivial [6]. In general, processes making use of channel mobility (delegation) and session interleaving, or dynamic network topologies with variable number of processes, are supported by our approach (within the limits imposed by the type systems), but are challenging to handle in top-down approaches. Inference of progress properties akin to lock freedom for session-based calculi has been studied in [17,3], although only finite types are considered in these works.

The reconstruction of *global* protocol descriptions from *local* session types has been studied in [15,16]. In this respect, our work fills the remaining gap and provides a reconstruction tool from processes to local session types. We plan to investigate the integration with [15,16] in future work.

Acknowledgments. The authors are grateful to the reviewers for their detailed comments and useful suggestions. The first two authors have been supported by Ateneo/CSP project SALT. The first author has also been supported by ICT COST Action IC1201 BETTY and MIUR project CINA.

#### References

- 1. T. Amtoft, F. Nielson, and H. Nielson. *Type and effect systems: behaviours for concurrency*. Imperial College Press, 1999.
- T. Amtoft, F. Nielson, and H. R. Nielson. Type and behaviour reconstruction for higher-order concurrent programs. J. Funct. Program., 7(3):321–347, 1997.
- M. Coppo, M. Dezani-Ciancaglini, L. Padovani, and N. Yoshida. Inference of global progress properties for dynamically interleaved multiparty sessions. In *COORDINATION'13*, LNCS 7890, pages 45–59. Springer, 2013.
- 4. B. Courcelle. Fundamental properties of infinite trees. Theor. Comp. Sci., 25:95–169, 1983.
- O. Dardha, E. Giachino, and D. Sangiorgi. Session types revisited. In *PPDP'12*, pages 139–150. ACM, 2012.
- P.-M. Deniélou and N. Yoshida. Multiparty session types meet communicating automata. In ESOP'12, LNCS 7211, pages 194–213. Springer, 2012.
- F. Henglein. Type inference with polymorphic recursion. ACM Trans. Program. Lang. Syst., 15(2):253–289, Apr. 1993.
- K. Honda, N. Yoshida, and M. Carbone. Multiparty asynchronous session types. In POPL'08, pages 273–284. ACM, 2008.
- 9. A. Igarashi and N. Kobayashi. Type reconstruction for linear  $\pi$ -calculus with I/O subtyping. *Inf. and Comp.*, 161(1):1–44, 2000.
- A. J. Kfoury, J. Tiuryn, and P. Urzyczyn. Type reconstruction in the presence of polymorphic recursion. ACM Trans. Program. Lang. Syst., 15(2):290–311, 1993.
- 11. N. Kobayashi. A type system for lock-free processes. Inf. and Comp., 177(2):122–159, 2002.
- N. Kobayashi. Type-based information flow analysis for the pi-calculus. *Acta Informatica*, 42(4-5):291–347, 2005.
- N. Kobayashi. A new type system for deadlock-free processes. In CONCUR'06, LNCS 4137, pages 233–247. Springer, 2006.
- N. Kobayashi, B. C. Pierce, and D. N. Turner. Linearity and the pi-calculus. ACM Trans. Program. Lang. Syst., 21(5):914–947, 1999.
- J. Lange and E. Tuosto. Synthesising choreographies from local session types. In CON-CUR'12, LNCS 7454, pages 225–239. Springer, 2012.
- J. Lange, E. Tuosto, and N. Yoshida. From communicating machines to graphical choreographies. In *POPL'15*, pages 221–232. ACM, 2015.
- 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.
- L. Padovani. Deadlock and lock freedom in the linear π-calculus. In CSL-LICS'14, pages 72:1-72:10. ACM, 2014. http://hal.archives-ouvertes.fr/hal-00932356v2/.
- 19. L. Padovani. Type reconstruction for the linear  $\pi$ -calculus with composite and equi-recursive types. In *FoSSaCS'14*, LNCS 8412, pages 88–102. Springer, 2014.