

Soft Session Types

Ugo Dal Lago

Università di Bologna
INRIA Sophia Antipolis
dallago@cs.unibo.it

Paolo Di Giamberardino

Laboratoire d'Informatique
de Paris Nord
digiambe@lipn.univ-paris13.fr

We show how systems of sessions types can enforce interactions to be bounded for all typable processes. The type system we propose is based on Lafont's soft linear logic and is strongly inspired by recent works about session types as intuitionistic linear logic formulas. Our main result is the existence, for every typable process, of a polynomial bound on the length of any reduction sequence starting from it and on the size of any of its reducts.

1 Introduction

Session types are one of the most successful paradigms around which communication can be disciplined in a concurrent or object-based environment. They can come in many different flavors, depending on the underlying programming language and on the degree of flexibility they allow when defining the structure of sessions. As an example, systems of session types for multi-party interaction have been recently introduced [9], while a form of higher-order session has been shown to be definable [12]. Recursive types, on the other hand, have been part of the standard toolset of session type theories since their inception [8].

The key property induced by systems of session types is the following: if two (or more) processes can be typed with “dual” session types, then they can interact with each other without “going wrong”, i.e. avoiding situations where one party needs some data with a certain type and the other(s) offer something of a different, incompatible type. Sometimes, one would like to go beyond that and design a type system which guarantees stronger properties, including quantitative ones. An example of a property that we find particularly interesting is the following: suppose that two processes P and Q interact by creating a session having type A through which they communicate. Is this interaction guaranteed to be finite? How long would it last? Moreover, P and Q could be forced to interact with other processes in order to be able to offer A . The question could then become: can the global amount of interaction be kept under control? In other words, one could be interested in *proving the interaction induced by sessions to be bounded*. This problem has been almost neglected by the research community in the area of session types, although it is the *manifesto* of the so-called implicit computational complexity (ICC), where one aims at giving machine-free characterizations of complexity classes based on programming languages and logical systems.

Linear logic (LL in the following) has been introduced twenty-five years ago by Jean-Yves Girard [7]. One of its greatest merits has been to allow a finer analysis of the computational content of both intuitionistic and classical logic. In turn, this is possible by distinguishing multiplicative as well as additive connectives, by an involutive notion of negation, and by giving a new status to structural rules allowing them to be applicable only to modal formulas. One of the many consequences of this new, refined way of looking at proof theory has been the introduction of natural characterizations of complexity classes by fragments of linear logic. This is possible because linear logic somehow “isolates” complexity in the modal fragment of the logic (which is solely responsible for the hyperexponential complexity of cut

elimination in, say intuitionistic logic), which can then be restricted so as to get exactly the expressive power needed to capture small complexity classes. One of the simplest and most elegant of those systems is Lafont’s soft linear logic (SLL in the following), which has been shown to correspond to polynomial time in the realm of classical [10], quantum [6] and higher-order concurrent computation [5].

Recently, Caires and Pfenning [1] have shown how a system of session types can be built around intuitionistic linear logic, by introducing π DILL, a type system for the π -calculus where types and rules are derived from the ones of intuitionistic linear logic. In their system, multiplicative connectives like \otimes and \multimap allow to model sequentiality in sessions, while the additive connectives $\&$ and \oplus model external and internal choice, respectively. The modal connective $!$, on the other hand, allows to model a server of type $!A$ which can offer the functionality expressed by A multiple times.

In this paper, we study a restriction of π DILL, called π DSLL, which can be thought of as being derived from π DILL in the same way as SLL is obtained from LL. In other words, the operator $!$ behaves in π DSLL in the same way as in SLL. The main result we prove about π DSLL is precisely about bounded interaction: whenever P can be typed in π DSLL and $P \rightarrow^n Q$, then both n and $|Q|$ (the size of the process Q , to be defined later) are polynomially related to $|P|$. This ensures an abstract but quite strong form of bounded interaction. Another, perhaps more “interactive” formulation of the same result is the following: if P and Q interact via a channel of type A , then the “complexity” of this interaction is bounded by a polynomial on $|P| + |Q|$, whose degree only depends on A .

We see this paper as the first successful attempt to bring techniques from implicit computational complexity into the realm of session types. Although proving bounded interaction has been technically nontrivial, due to the peculiarities of the π -calculus, we think the main contribution of this work lies in showing that bounded termination can be enforced by a natural adaptation of known systems of session types. An extended version with more details is available [4].

2 π DILL, an Informal Account

In this section, we will outline the main properties of π DILL, a session type system recently introduced by Caires and Pfenning [1, 2]. For more information, please consult the two cited papers.

In π DILL, session types are nothing more than formulas of (propositional) intuitionistic linear logic without atoms but with (multiplicative) constants:

$$A ::= \mathbf{1} \mid A \otimes A \mid A \multimap A \mid A \oplus A \mid A \& A \mid !A.$$

These types are assigned to channels (names) by a formal system deriving judgments in the form

$$\Gamma; \Delta \vdash P :: x : A,$$

where Γ and Δ are contexts assigning types to channels, and P is a process of the name-passing π calculus. The judgment above can be read as follows: the process P acts on the channel x according to the session type A *whenever* composed with processes behaving according to Γ and Δ (each on a specific channel). Informally, the various constructions on session types can be explained as follows:

- $\mathbf{1}$ is the type of an empty session channel. A process offering to communicate via a session channel typed this way simply synchronizes with another process through it without exchanging anything. This is meant to be an abstraction for all ground session types, e.g. natural numbers, lists, etc. In linear logic, this is the unit for \otimes .

- $A \otimes B$ is the type of a session channel x through which a message carrying another channel with type A is sent. After performing this action, the underlying process behaves according to B on the *same* channel x .
- $A \multimap B$ is the adjoint to $A \otimes B$: on a channel with this type, a process communicate by first performing an input and receiving a channel with type A , then acting according to B , again on x .
- $A \oplus B$ is the type of a channel on which a process either sends a special message `inl` and performs according to A or sends a special message `inr` and performs according to B .
- The type $A \& B$ can be assigned to a channel x on which the underlying process offers the possibility of choosing between proceeding according to A or to B , both on x . So, in a sense, $\&$ models external choice.
- Finally, the type $!A$ is attributed to a channel x only if a process can be replicated by receiving a channel y through x , then behaving on y according to A .

The assignments in Γ and Δ are of two different natures:

- An assignment of a type A to a channel x in Δ signals the need by P of a process offering a session of type A on the channel x ; for this reason, Δ is called the *linear context*;
- An assignment of a type A to a channel x in Γ , on the other hand, represents the need by P of a process offering a session of type $!A$ on the channel x ; thus, Γ is the *exponential context*.

Typing rules π DILL are very similar to the ones of DILL, itself one of the many possible formulations of linear logic as a sequent calculus. In particular, there are two cut rules, each corresponding to a different portion of the context:

$$\frac{\Gamma; \Delta_1 \vdash P :: x : A \quad \Gamma; \Delta_2, x : A \vdash Q :: T}{\Gamma; \Delta_1, \Delta_2 \vdash (vx)(P \mid Q) :: T} \quad \frac{\Gamma; \emptyset \vdash P :: y : A \quad \Gamma, x : A; \Delta \vdash Q :: T}{\Gamma; \Delta \vdash (vx)(!x(y).P \mid Q) :: T}$$

Please observe how cutting a process P against an assumption in the exponential context requires to “wrap” P inside a replicated input: this allows to *turn P into a server*.

In order to illustrate the intuitions above, we now give an example. Suppose that a process P models a service which acts on x as follows: it receives two natural numbers, to be interpreted as the number and secret code of a credit card and, if they correspond to a valid account, returns an MP3 file and a receipt code to the client. Otherwise, the session terminates. To do so, P needs to interact with another service (e.g. a banking service) Q through a channel y . The banking service, among others, provides a way to verify whether a given number and code correspond to a valid credit card. In π DILL, the process P would receive the type

$$\emptyset; y : (\mathbf{N} \multimap \mathbf{1} \oplus \mathbf{1}) \& A \vdash P :: x : \mathbf{N} \multimap \mathbf{N} \multimap (\mathbf{S} \otimes \mathbf{N}) \oplus \mathbf{1},$$

where \mathbf{N} and \mathbf{S} are pseudo-types for natural numbers and MP3s, respectively. A is the type of all the other functionalities Q provides. As an example, P could be the following process:

$$\begin{aligned} & x(nm_1).x(cd_1).y.inl; \\ & (vnm_2)y\langle nm_2 \rangle.(vcd_2)y\langle cd_2 \rangle. \\ & y.case(x.inl; (vmp)x\langle mp \rangle.(vrp)x\langle rp \rangle, x.inr; 0) \end{aligned}$$

Observe how the credit card number and secret code forwarded to Q are not the ones sent by the client: the flow of information happening inside a process is abstracted away in π DILL. Similarly, one can write a process Q and assign it a type as follows: $\emptyset; \emptyset \vdash Q :: y : (\mathbf{N} \multimap \mathbf{1} \oplus \mathbf{1}) \& A$. Putting the two derivations together, we obtain $\emptyset; \emptyset \vdash (vx)(P \mid Q) :: x : \mathbf{N} \multimap \mathbf{N} \multimap (\mathbf{S} \otimes \mathbf{N}) \oplus \mathbf{1}$.

Let us now make an observation which will probably be appreciated by the reader familiar with linear logic. The processes P and Q can be typed in π DILL without the use of any exponential rule, nor

of cut. What allows to type the parallel composition $(\nu x)(P \mid Q)$, on the other hand, is precisely the cut rule. The interaction between P and Q corresponds to the elimination of that cut. Since there isn't any exponential around, this process must be finite, since the size of the underlying process shrinks at every single reduction step. From a process-algebraic point of view, on the other hand, the finiteness of the interaction is an immediate consequence of the absence of any replication in P and Q .

The banking service Q can only serve one single session and would vanish at the end of it. To make it into a *persistent server* offering the same kind of session to possibly many different clients, Q must be put into a replication, obtaining $R = !z(y).Q$. In R , the channel z can be given type $!((\mathbf{N} \multimap \mathbf{1} \oplus \mathbf{1}) \& A)$ in the empty context. The process P should be somehow adapted to be able to interact with R : before performing the two outputs on y , it's necessary to "spawn" R by performing an output on z and passing y to it. This way we obtain a process S such that

$$\emptyset; z : !((\mathbf{N} \multimap \mathbf{1} \oplus \mathbf{1}) \& A) \vdash S :: x : \mathbf{N} \multimap \mathbf{N} \multimap (\mathbf{S} \otimes \mathbf{N}) \oplus \mathbf{1},$$

and the composition $(\nu z)(S \mid R)$ can be given the same type as $(\nu x)(P \mid Q)$. Of course, S could have used the channel z more than once, initiating different sessions. This is meant to model a situation in which the same client interacts with the same server by creating more than one session with the same type, itself done by performing *more than one output* on the same channel. Of course, servers can themselves depend on other servers. And these dependencies are naturally modeled by the exponential modality of linear logic.

3 On Bounded Interaction

In π DILL, the possibility of modeling persistent servers which in turn depend on other servers makes it possible to type processes which exhibit a very complex and combinatorially heavy interactive behavior.

Consider the following processes, the first one parameterized on any $i \in \mathbb{N}$:

$$\begin{aligned} \text{dupser}_i &\doteq !x_i(y).(\nu z)x_{i+1}\langle z \rangle.(\nu w)x_{i+1}\langle w \rangle.; \\ \text{dupclient} &\doteq (\nu y)x_0\langle y \rangle. \end{aligned}$$

In π DILL, these processes can be typed as follows:

$$\begin{aligned} \emptyset; x_{i+1} : !\mathbf{1} \vdash \text{dupser}_i :: x_i : !\mathbf{1}; \\ \emptyset; x_0 : !\mathbf{1} \vdash \text{dupclient} :: z : \mathbf{1}. \end{aligned}$$

Then, for every $n \in \mathbb{N}$ one can type the parallel composition $\text{mulser}_{n+1} \doteq (\nu x_1 \dots x_n)(\text{dupser}_n \parallel \dots \parallel \text{dupser}_0)$ as follows

$$\emptyset; x_n : !\mathbf{1} \vdash \text{mulser}_n :: x_0 : !\mathbf{1}.$$

Informally, mulser_n is a persistent server which offers a session type $\mathbf{1}$ on a channel x_0 , provided a server with the same functionality is available on x_n . The process mulser_n is the parallel composition of n servers in the form dupser_i , each spawning two different sessions provided by dupser_{i+1} on the same channel x_{i+1} .

The process mulser_n cannot be further reduced. But notice that, once mulser_n and dupclient are

composed, the following exponential blowup is bound to happen:

$$\begin{aligned}
(\nu x_0)(mulser_n \mid dupclient) &\equiv (\nu x_0 \dots x_n)(dupser_n \parallel \dots \parallel dupser_0 \mid dupclient) \\
&\rightarrow (\nu x_0 \dots x_n)(dupser_n \parallel \dots \parallel dupser_1 \mid P_1) \\
&\rightarrow^2 (\nu x_1 \dots x_n)(dupser_n \parallel \dots \parallel dupser_2 \mid P_2 \mid P_2) \\
&\rightarrow^4 (\nu x_2 \dots x_n)(dupser_n \parallel \dots \parallel dupser_3 \mid \underbrace{P_3 \parallel \dots \parallel P_3}_{4 \text{ times}}) \\
&\rightarrow^* (\nu x_n)(dupser_n \mid \underbrace{P_n \parallel \dots \parallel P_n}_{2^n \text{ times}}) \\
&\rightarrow^{2^n} 0.
\end{aligned}$$

Here, for every $i \in \mathbb{N}$ the process P_i is simply $(\nu y)x_i \langle y \rangle . (\nu z)x_i \langle z \rangle$. Notice that *both* the number of reduction steps *and* the size of intermediate processes are exponential in n , while the size of the initial process is linear in n . This is a perfectly legal process in π DILL. Moreover the type $\mathbf{!1}$ of the channel x_0 through which *dupclient* and *mulser_n* communicate does not contain any information about the “complexity” of the interaction: it is the same for every n .

The deep reasons why this phenomenon can happen lie in the very general (and “generous”) rules governing the behavior of the exponential modality $\mathbf{!}$ in linear logic. It is this generality that allows the embedding of propositional intuitionistic logic into linear logic. Since the complexity of normalization for the former [13, 11] is nonelementary, the exponential blowup described above is not a surprise.

It would be desirable, on the other hand, to be sure that the interaction caused by any process P is bounded: whenever $P \rightarrow^n Q$, then there’s a *reasonably low* upper bound to both n and $|Q|$. This is precisely what we achieve by restricting π DILL into π DSLlL.

4 π DSLlL: Syntax and Main Properties

In this section, the syntax of π DSLlL will be introduced. Moreover, some basic operational properties will be given.

4.1 The Process Algebra

π DSLlL is a type system for a fairly standard π -calculus, exactly the one on top of which π DILL is defined:

Definition 1 (Processes) *Given an infinite set of names or channels x, y, z, \dots , the set of processes is defined as follows:*

$$P ::= 0 \mid P \mid Q \mid (\nu x)P \mid x(y).P \mid x \langle y \rangle . P \mid \mathbf{!}x(y).P \mid x.\text{inl};P \mid x.\text{inr};P \mid x.\text{case}(P, Q)$$

The only non-standard constructs are the last three, which allow to define a choice mechanism: the process $x.\text{case}(P, Q)$ can evolve as P or as Q *after* having received a signal in the form inl or inr through x . Processes sending such a signal through the channel x , then continuing like P are, respectively, $x.\text{inl};P$ and $x.\text{inr};P$. The set of names occurring free in the process P (hereby denoted $fn(P)$) is defined as usual. The same holds for the capture avoiding substitution of a name x for y in a process P (denoted $P\{x/y\}$), and for α -equivalence between processes (denoted \equiv_α).

Structural congruence is an equivalence relation identifying those processes which are syntactically different but can be considered equal for very simple structural reasons:

Definition 2 (Structural Congruence) *The relation \equiv , called structural congruence, is the least congruence on processes satisfying the following seven axioms:*

$$\begin{aligned}
P &\equiv Q \quad \text{whenever } P \equiv_{\alpha} Q; & (\nu x)0 &\equiv 0; \\
P \mid 0 &\equiv P; & (\nu x)(\nu y)P &\equiv (\nu y)(\nu x)P; \\
P \mid Q &\equiv Q \mid P; & ((\nu x)P) \mid Q &\equiv (\nu x)(P \mid Q) \quad \text{whenever } x \notin \text{fn}(Q); \\
P \mid (Q \mid R) &\equiv (P \mid Q) \mid R.
\end{aligned}$$

Formal systems for reduction and labelled semantics can be defined in a standard way. We refer the reader to [1] for more details.

A quantitative attribute of processes which is delicate to model in process algebras is their *size*: how can we measure the size of a process? In particular, it is not straightforward to define a measure which both reflects the “number of symbols” in the process and is invariant under structural congruence (this way facilitating all proofs). A good compromise is the following:

Definition 3 (Process Size) *The size $|P|$ of a process P is defined by induction on the structure of P as follows:*

$$\begin{aligned}
|0| &= 0; & |x(y).P| &= |P| + 1; & |x.\text{inl};P| &= |P| + 1; \\
|P \mid Q| &= |P| + |Q|; & |x\langle y \rangle.P| &= |P| + 1; & |x.\text{inr};P| &= |P| + 1; \\
|(\nu x)P| &= |P|; & |!x(y).P| &= |P| + 1; & |x.\text{case}(P, Q)| &= |P| + |Q| + 1.
\end{aligned}$$

According to the definition above, the empty process 0 has null size, while restriction does not increase the size of the underlying process. This allows for a definition of size which remains invariant under structural congruence. The price to pay is the following: the “number of symbols” of a process P can be arbitrarily bigger than $|P|$ (e.g. for every $n \in \mathbb{N}$, $|(\nu x)^n P| = |P|$). However, we have the following:

Lemma 1 *For every P, Q , $|P| = |Q|$ whenever $P \equiv Q$. Moreover, there is a polynomial p such that for every P , there is Q with $P \equiv Q$ and the number of symbols in Q is at most $p(|Q|)$.*

4.2 The Type System

The language of types of πDSLL is exactly the same as the one of πDILL , and the interpretation of type constructs does not change (see Section 2 for some informal details). Typing judgments and typing rules, however, are significantly different, in particular, in the treatment of the exponential connective $!$.

Typing judgments become syntactical expressions in the form

$$\Gamma; \Delta; \Theta \vdash P :: x : A.$$

First of all, observe how the context is divided into *three* chunks now: Γ and Δ have to be interpreted as exponential contexts, while Θ is the usual linear context from πDILL . The necessity of having *two* exponential contexts is a consequence of the finer, less canonical exponential discipline of SLL compared to the one of LL. We use the following terminology: Γ is said to be the *auxiliary* context, while Δ is the *multiplexor* context.

Typing rules are in Figure 1. The rules governing the typing constant $\mathbf{1}$, the multiplicatives (\otimes and \multimap) and the additives (\oplus and $\&$) are exact analogues of the ones from πDILL . The only differences come from the presence of two exponential contexts: in binary multiplicative rules ($\otimes\text{R}$ and $\multimap\text{L}$) the auxiliary

$$\begin{array}{c}
\frac{\Gamma; \Delta; \Theta \vdash P :: T}{\Gamma; \Delta; \Theta, x : \mathbf{1} \vdash P :: T} \mathbf{1L} \quad \frac{}{\Gamma; \Delta; \emptyset \vdash 0 :: x : \mathbf{1}} \mathbf{1R} \\
\\
\frac{\Gamma; \Delta; \Theta, y : A, x : B \vdash P :: T}{\Gamma; \Delta; \Theta, x : A \otimes B \vdash x(y).P :: T} \otimes L \quad \frac{\Gamma_1; \Delta; \Theta_1 \vdash P :: y : A \quad \Gamma_2; \Delta; \Theta_2 \vdash Q :: x : B}{\Gamma_1, \Gamma_2; \Delta; \Theta_1, \Theta_2 \vdash (vy)x(y).(P | Q) :: x : A \otimes B} \otimes R \\
\\
\frac{\Gamma_1; \Delta; \Theta_1, y : A \vdash P :: T \quad \Gamma_2; \Delta; \Theta_2, x : B \vdash Q :: T}{\Gamma_1, \Gamma_2; \Delta; \Theta_1, \Theta_2, x : A \multimap B \vdash (vy)x(y).(P | Q) :: T} \multimap L \quad \frac{\Gamma; \Delta; \Theta, y : A \vdash P :: x : B}{\Gamma; \Delta; \Theta \vdash x(y).P :: x : A \multimap B} \multimap R \\
\\
\frac{\Gamma; \Delta; \Theta, x : A \vdash P :: T \quad \Gamma; \Delta; \Theta, x : B \vdash P :: T}{\Gamma; \Delta; x : A \oplus B, \Theta \vdash y.\text{case}(P, Q) :: T} \oplus L \quad \frac{\Gamma; \Delta; \Theta \vdash P :: x : A}{\Gamma; \Delta; \Theta \vdash x.\text{inl}; P :: x : A \oplus B} \oplus R_1 \\
\\
\frac{\Gamma; \Delta; \Theta \vdash P :: x : B}{\Gamma; \Delta; \Theta \vdash x.\text{inr}; P :: x : A \oplus B} \oplus R_2 \quad \frac{\Gamma; \Delta; \Theta, x : A \vdash P :: T}{\Gamma; \Delta; \Theta, x : A \& B \vdash x.\text{inl}; P :: T} \& L_1 \\
\\
\frac{\Gamma; \Delta; \Theta, x : B \vdash P :: T}{\Gamma; \Delta; \Theta, x : A \& B \vdash x.\text{inr}; P :: T} \& L_2 \quad \frac{\Gamma; \Delta; \Theta \vdash P :: x : A \quad \Gamma; \Delta; \Theta \vdash P :: x : B}{\Gamma; \Delta; \Theta \vdash y.\text{case}(P, Q) :: x : A \& B} \& R \\
\\
\frac{\Gamma; \Delta, x : A; \Theta, y : A \vdash P :: T}{\Gamma; \Delta, x : A; \Theta \vdash (vy)x(y).P :: T} b_{\#} \quad \frac{\Gamma; \Delta; \Theta, y : A \vdash P :: T}{\Gamma, x : A; \Delta; \Theta \vdash (vy)x(y).P :: T} b_! \\
\\
\frac{\Gamma; \Delta, x : A; \Theta \vdash P :: T}{\Gamma; \Delta; \Theta, x : !A \vdash P :: T} !L_{\#} \quad \frac{\Gamma, x : A; \Delta; \Theta \vdash P :: T}{\Gamma; \Delta; \Theta, x : !A \vdash P :: T} !L_! \quad \frac{\Gamma; \emptyset; \emptyset \vdash Q :: y : A}{\emptyset; \Delta; !\Gamma \vdash !x(y).Q :: x : !A} !R \\
\\
\frac{\Gamma_1; \Delta; \Theta_1 \vdash P :: x : A \quad \Gamma_2; \Delta; \Theta_2, x : A \vdash Q :: T}{\Gamma_1, \Gamma_2; \Delta; \Theta_1, \Theta_2 \vdash (vx)(P | Q) :: T} \text{cut} \quad \frac{\Delta; \emptyset; \emptyset \vdash P :: y : A \quad \Gamma; \Delta, x : A; \Theta \vdash Q :: T}{\Gamma; \Delta; \Theta \vdash (vx)(!x(y).P | Q) :: T} \text{cut}_{\#} \\
\\
\frac{\Gamma_1; \emptyset; \emptyset \vdash P :: y : A \quad \Gamma_2, x : A; \Delta; \Theta \vdash Q :: T}{\Gamma_1, \Gamma_2; \Delta; \Theta \vdash (vx)(!x(y).P | Q) :: T} \text{cut}_!
\end{array}$$

Figure 1: Typing rules for πDSLL .

context is treated multiplicatively, while the multiplexor context is treated additively, as in πDILL ¹. Now, consider the rules governing the exponential connective $!$, which are $b_!$, $b_{\#}$, $!L_!$, $!L_{\#}$ and $!R$:

- The rules $b_!$ and $b_{\#}$ both allow to spawn a server. This corresponds to turning an assumption $x : A$ in the linear context into one $y : A$ in one of the exponential contexts; in $b_{\#}$, $x : A$ could be already present in the multiplexor context, while in $b_!$ this cannot happen;
- The rules $!L_!$ and $!L_{\#}$ lift an assumption in the exponential contexts to the linear context; this requires changing its type from A to $!A$;
- The rule $!R$ allows to turn an ordinary process into a server, by packaging it into a replicated input and modifying its type.

¹The reader familiar with linear logic and proof nets will recognize in the different treatment of the auxiliary and multiplexor contexts, one of the basic principles of SLL: *contraction is forbidden on the auxiliary doors of exponential boxes*. The channel names contained in the auxiliary context correspond to the auxiliary doors of exponential boxes, so we treat them multiplicatively. The contraction effect induced by the additive treatment of the channel names in the multiplexor context corresponds to the multiplexing rule of SLL.

Finally there are *three* cut rules in the system, namely cut , $\text{cut}_!$ and $\text{cut}_\#$:

- cut is the usual linear cut rule, i.e. the natural generalization of the one from πDILL .
- $\text{cut}_!$ and $\text{cut}_\#$ allow to eliminate an assumption in one of the the two exponential contexts. In both cases, the process which allows to do that must be typable with empty linear and multiplexor contexts.

4.3 Back to Our Example

Let us now reconsider the example processes introduced in Section 3. The basic building block over which everything is built was the process $\text{dupser}_i = !x_i(y).(\nu z)x_{i+1}\langle z \rangle.(\nu w)x_{i+1}\langle w \rangle$. We claim that for every i , the process dupser_i is *not* typable in πDSSL . To understand why, observe that the only way to type a replicated input like dupser_i is by the typing rule $!R$, and that its premise requires the body of the replicated input to be typable with empty linear and multiplexor contexts. A quick inspection on the typing rules reveals that every name in the *auxiliary* context occurs (free) exactly once in the underlying process (provided we count two occurrences in the branches of a case as just *a single* occurrence). However, the name x_{i+1} appears *twice* in the body of dupser_i .

A slight variation on the example above, on the other hand, *can* be typed in πDSSL , but this requires changing its type. See [4] for more details.

4.4 Subject Reduction

A basic property most type systems for functional languages satisfy is subject reduction: typing is preserved along reduction. For processes, this is often true for internal reduction: if $P \rightarrow Q$ and $\vdash P : A$, then $\vdash Q : A$. In this section, a subject reduction result for πDSSL will be given and some ideas on the underlying proof will be described. Some concepts outlined here will become necessary ingredients in the proof of bounded interaction, to be done in Section 5 below. Subject reduction is proved by closely following the path traced by Caires and Pfenning; as a consequence, we proceed quite quickly, concentrating our attention on the differences with their proof.

When proving subject reduction, one constantly work with type derivations. This is particularly true here, where (internal) reduction corresponds to the cut-elimination process. A linear notation for proofs in the form of *proof terms* can be easily defined, allowing for more compact descriptions. As an example, a proof in the form

$$\frac{\pi : \Gamma_1; \Delta; \Theta_1 \vdash P :: x : A \quad \rho : \Gamma_2; \Delta; \Theta_2, x : A \vdash Q :: T}{\Gamma_1, \Gamma_2; \Delta; \Theta_1, \Theta_2 \vdash (\nu x)(P \mid Q) :: T} \text{ cut}$$

corresponds to the proof term $\text{cut}(D, x.E)$, where D is the proof term for π and E is the proof term for ρ . If D is a proof term corresponding to a type derivation for the process P , we write $\widehat{D} = P$. From now on, proof terms will often take the place of processes: $\Gamma; \Delta; \Theta \vdash D :: T$ stands for the existence of a type derivation D with conclusion $\Gamma; \Delta; \Theta \vdash \widehat{D} :: T$. A proof term D is said to be normal if it does not contain any instances of cut rules.

Subject reduction will be proved by showing that if P is typable by a type derivation D and $P \rightarrow Q$, then a type derivation E for Q exists. Actually, E can be obtained by manipulating D using techniques derived from cut-elimination. Noticeably, not every cut-elimination rule is necessary to prove subject reduction. In other words, we are in presence of a weak correspondence between proof terms and processes, and remain far from a genuine Curry-Howard correspondence.

Those manipulations of proof-terms which are necessary to prove subject reduction can be classified as follows:

- First of all, a binary relation \Longrightarrow on proof terms called *computational reduction* can be defined. At the logical level, this corresponds to proper cut-elimination steps, i.e. those cut-elimination steps in which two rules introducing the same connective interact. At the process level, computational reduction correspond to internal reduction. \Longrightarrow is not symmetric.
- A binary relation \mapsto on proof terms called *shift reduction*, distinct from \Longrightarrow must be introduced. At the process level, it corresponds to structural congruence. As \Longrightarrow , \mapsto is not a symmetric relation.
- Finally, an equivalence relation \equiv on proof terms called *proof equivalence* is necessary. At the logical level, this corresponds to the so-called commuting conversions, while at the process level, the induced processes are either structurally congruent or strongly bisimilar.

The reflexive and transitive closure of $\mapsto \cup \equiv$ is denoted with \leftrightarrow , i.e. $\leftrightarrow = (\mapsto \cup \equiv)^*$. There is not enough space here to give the rules defining \Longrightarrow , \mapsto and \equiv . Let us give only some relevant examples:

- Let us consider the proof term $D = \text{cut}((\otimes R(F, G)), x. \otimes L(x, y.x.H))$ which corresponds to the \otimes -case of cut elimination. By a computational reduction rule, $D \Longrightarrow E = \text{cut}(F, y.\text{cut}(G, x.H))$. From the process side, $\widehat{D} = (vx)((vy)x(y).\widehat{F} \mid \widehat{G}) \mid x(y).\widehat{H}$ and $\widehat{E} = (vx)(vy)((\widehat{F} \mid \widehat{G}) \mid \widehat{H})$, where \widehat{E} is the process obtained from \widehat{D} by internal passing the channel y through the channel x .
- Let $D = \text{cut}(!R(F, x_1, \dots, x_n), x. !L_!(x.G))$ be the proof obtained by composing a proof F (whose last rule is $!R$) with a proof G (whose last rule is $!L_!$) through a cut rule. A shift reduction rule tells us that $D \mapsto E = !L_!(x_1. !L_!(x_2. \dots !L_!(x_n. \text{cut}_!(F, y.G)) \dots))$, which corresponds to the opening of a box in SLL. The shift reduction does not have a corresponding reduction step at process level, since $\widehat{D} \equiv \widehat{E}$; nevertheless, it is defined as an asymmetric relation, for technical reasons connected to the proof of bounded interaction.
- Let $D = \text{cut}_\#(F, x.\text{cut}(G, y.H))$. A defining rule for proof equivalence \equiv , states that in D the $\text{cut}_\#$ rule can be permuted over the cut rule, by duplicating F ; namely $D \equiv E = \text{cut}(\text{cut}_\#(F, x.G), y.\text{cut}_\#(F, x.H))$. This is possible because the channel x belongs to the multiplexor contexts of both G, H , such contexts being treated additively. At the process level, $\widehat{D} = (vx)((!x(y).\widehat{F}) \mid (vy)(\widehat{G} \mid \widehat{H}))$, while $\widehat{E} = (vy)((vx)(!x(y).\widehat{F}) \mid \widehat{G}) \mid ((vx)(!x(y).\widehat{F}) \mid \widehat{H})$, \widehat{D} and \widehat{E} being strongly bisimilar.

Before proceeding to Subject Reduction, we give the following two lemmas, concerning structural properties of the type system:

Lemma 2 (Weakening lemma) *If $\Gamma; \Delta; \Theta \vdash D :: T$ and whenever $\Delta \subseteq \Phi$, it holds that $\Gamma; \Phi; \Theta \vdash D :: T$.*

Proof. By a simple induction on the structure of D . □

Lemma 3 (Lifting lemma) *If $\Gamma; \Delta; \Theta \vdash D :: T$ then there exists an E such that $\emptyset; \Gamma, \Delta; \Theta \vdash E :: T$ where $\widehat{E} = \widehat{D}$. We denote E by D_{\downarrow} .*

Proof. Again, a simple induction on the structure of the proof term D . □

Finally:

Theorem 1 (Subject Reduction) *Let $\Gamma; \Delta; \Theta \vdash D :: T$. Suppose that $\widehat{D} = P \rightarrow Q$. Then there is E such that $\widehat{E} = Q$, $D \leftrightarrow \Longrightarrow \leftrightarrow E$ and $\Phi; \Psi; \Theta \vdash E :: T$, where $\Gamma, \Delta = \Phi, \Psi$.*

Let us give a sketch of the proof of Theorem 1. We reason by induction on the structure of D . Since $\widehat{D} = P \rightarrow Q$ the only possible last rules of D can be: $\mathbf{1L}$, $!L_!$, $!L_\#$, a linear cut (cut) or an exponential cut ($\text{cut}_!$ or $\text{cut}_\#$). In all the other cases, the underlying process can only perform a visible action, as can be easily verified by inspecting the rules from Figure 1. With this observation in mind, let us inspect the operational semantics derivation proving that $P \rightarrow Q$. At some point we will find two subprocesses of P ,

call them R and S , which communicate, causing an internal reduction. We here claim that this can only happen in presence of a cut, and only the communication between R and S must occur along the channel involved in the cut. Now, it's only a matter of showing that the just described situation can be “resolved” preserving types. And this can be done by way of several lemmas, like the following:

Lemma 4 *Assume that:*

1. $\Gamma_1; \Delta; \Theta_1 \vdash D :: x : A \otimes B$ with $\widehat{D} = P \xrightarrow{(\nu y)x(y)} Q$;
2. $\Gamma_2; \Delta; \Theta_2, x : A \otimes B \vdash E :: z : C$ with $\widehat{E} = R \xrightarrow{x(y)} S$.

Then:

1. $\text{cut}(D, x.E) \leftrightarrow \implies \hookrightarrow F$ for some F ;
2. $\Gamma_1, \Gamma_2; \Delta; \Theta_1, \Theta_2 \vdash F :: z : C$, where $\widehat{F} \equiv (\nu x)(Q \mid S)$.

The other lemmas can be found in [4]. By the way, this proof technique is very similar to the one introduced by Caires and Pfenning [1].

5 Proving Polynomial Bounds

In this section, we prove the main result of this paper, namely some polynomial bounds on the length of internal reduction sequences and on the size of intermediate results for processes typable in πDILL . In other words, interaction will be shown to be bounded. The simplest formulation of this result is the following:

Theorem 2 *For every type A , there is a polynomial p_A such that whenever $\emptyset; \emptyset; x : A \vdash D :: y : \mathbf{I}$ and $\emptyset; \emptyset; \emptyset \vdash E :: x : A$ where D and E are normal and $(\nu x)(\widehat{D} \mid \widehat{E}) \rightarrow^n P$, it holds that $n, |P| \leq p_A(|\widehat{D}| + |\widehat{E}|)$*

Intuitively, what Theorem 2 says is that the complexity of the interaction between two processes typable without cuts and communicating through a channel with session type A is polynomial in their sizes, where the specific polynomial involved only depends on A itself. In other words, the complexity of the interaction is not only bounded, but can be somehow “read off” from the types of the communicating parties.

How does the proof of Theorem 2 look like? Conceptually, it can be thought of as being structured into four steps:

1. First of all, a natural number $\mathbb{W}(D)$ is attributed to any proof term D . $\mathbb{W}(D)$ is said to be the *weight* of D .
2. Secondly, the weight of any proof term is shown to strictly decrease along computational reduction, not to increase along shifting reduction and to stay the same for equivalent proof terms.
3. Thirdly, $\mathbb{W}(D)$ is shown to be bounded by a polynomial on $|\widehat{D}|$, where the exponent only depends on the nesting depth of boxes of D , denoted $\mathbb{B}(D)$.
4. Finally, the box depth $\mathbb{B}(D)$ of any proof term D is shown to be “readable” from its type interface.

This is exactly what we are going to do in the rest of this section. Please observe how points 1–3 above allow to prove the following stronger result, from which Theorem 2 easily follows, given point 4:

Proposition 1 *For every $n \in \mathbb{N}$, there is a polynomial p_n such that for every process P with $\Gamma; \Delta; \Theta \vdash P :: T$, if $P \rightarrow^m Q$, then $m, |Q| \leq p_{\mathbb{B}(P)}(|P|)$.*

5.1 Preliminary Definitions

Some concepts have to be given before we can embark in the proof of Proposition 1. First of all, we need to define what the box-depth of a process is. Simply, given a process P , its *box-depth* $\mathbb{B}(P)$ is the

nesting-level of replications² in P . As an example, the box-depth of $!x(y).!z(w).0$ is 2, while the one of $(\nu x)y(z)$ is 0. Analogously, the box-depth of a proof term D is simply $\mathbb{B}(\widehat{D})$.

Now, suppose that $\Gamma; \Delta; \Theta \vdash D :: T$ and that $x : A$ belongs to either Γ or Δ , i.e. that x is an “exponential” channel in D . A key parameter is the *virtual number of occurrences* of x in D , which is denoted as $\mathbb{FO}(x, D)$. This parameter, as its name suggests, is not simply the number of literal occurrences of x in D , but takes into account possible duplications derived from cuts. So, for example, $\mathbb{FO}(w, \text{cut}_!(D, x.E)) = \mathbb{FO}(x, E) \cdot \mathbb{FO}(w, D) + \mathbb{FO}(w, E)$, while $\mathbb{FO}(w, \otimes R(D, E))$ is merely $\mathbb{FO}(w, D) + \mathbb{FO}(w, E)$. Obviously, $\mathbb{FO}(w, b_!(x, w.D)) = 1$ and $\mathbb{FO}(w, b_\#(x, w.D)) = 1$.

A channel in either the auxiliary or the exponential context can “float” to the linear context as an effect of rules $!L_!$ or $!L_\#$. From that moment on, it can only be treated as a linear channel. As a consequence, it makes sense to define the *duplicability factor* of a proof term D , written $\mathbb{D}(D)$, simply as the maximum of $\mathbb{FO}(x, D)$ over all instances of the rules $!L_!$ or $!L_\#$ in D , where x is the involved channel. For example, $\mathbb{D}(!L_!(x.D)) = \max\{\mathbb{D}(D), \mathbb{FO}(y, D)\}$ and $\mathbb{D}(\multimap L(x, D, y.E)) = \max\{\mathbb{D}(D), \mathbb{D}(E)\}$.

It’s now possible to give the definition of $\mathbb{W}(D)$, namely the *weight* of the proof term D . Before doing that, however, it is necessary to give a parameterized notion of weight, denoted $\mathbb{W}_n(D)$. Intuitively, $\mathbb{W}_n(D)$ is defined similarly to $|\widehat{D}|$. However, every input and output action in \widehat{D} can possibly count for more than one:

- Everything inside D in $!R(x_1, \dots, x_n, D)$ counts for n ;
- Everything inside D in either $\text{cut}_!(D, x.E)$ or $\text{cut}_\#(D, x.E)$ counts for $\mathbb{FO}(x, E)$.

For example, $\mathbb{W}_n(\text{cut}_\#(D, x.E)) = \mathbb{FO}(x, E) \cdot \mathbb{W}_n(D) + \mathbb{W}_n(E)$, while $\mathbb{W}_n(\&L_2(x, y.D)) = 1 + \mathbb{W}_n(D)$. Now, $\mathbb{W}(D)$ is simply $\mathbb{W}_{\mathbb{D}(D)}(D)$. The concepts we have just introduced are more precisely defined in [4].

5.2 Monotonicity Results

The crucial ingredient for proving polynomial bounds are a series of results about how the weight D evolves when D is put in relation with another proof term E by way of either \Longrightarrow , \mapsto or \equiv . Whenever a proof term D computationally reduces to E , the underlying weight is guaranteed to strictly decrease:

Proposition 2 *If $\Gamma; \Delta; \Theta \vdash D :: T$ and $D \Longrightarrow E$, then $\Phi; \Psi; \Theta \vdash E :: T$ (where $\Gamma, \Delta = \Phi, \Psi$), $\mathbb{D}(E) \leq \mathbb{D}(D)$ and $\mathbb{W}(E) < \mathbb{W}(D)$.*

Proof. By induction on the proof that $D \Longrightarrow E$. Some interesting cases:

- Suppose that $D = \text{cut}(\multimap R(y.F), x. \multimap L(x, G, x.H)) \Longrightarrow \text{cut}(\text{cut}(G, y.F), x.H) = E$. Then,

$$\begin{aligned} \mathbb{D}(D) &= \max\{\mathbb{D}(F), \mathbb{D}(G), \mathbb{D}(H)\} = \mathbb{D}(E); \\ \mathbb{W}(D) &= \mathbb{W}_{\mathbb{D}(D)}(D) = 3 + \mathbb{W}_{\mathbb{D}(D)}(F) + \mathbb{W}_{\mathbb{D}(D)}(G) + \mathbb{W}_{\mathbb{D}(D)}(H) \\ &> 2 + \mathbb{W}_{\mathbb{D}(E)}(F) + \mathbb{W}_{\mathbb{D}(E)}(G) + \mathbb{W}_{\mathbb{D}(E)}(H) = \mathbb{W}_{\mathbb{D}(E)}(E) = \mathbb{W}(E). \end{aligned}$$

- Suppose that $D = \text{cut}(\&R(F, G), x. \&L_1(x, y.H)) \Longrightarrow \text{cut}(F, x.H) = E$. Then,

$$\begin{aligned} \mathbb{D}(D) &= \max\{\mathbb{D}(F), \mathbb{D}(G), \mathbb{D}(H)\} = \mathbb{D}(E); \\ \mathbb{W}(D) &= \mathbb{W}_{\mathbb{D}(D)}(D) = 3 + \mathbb{W}_{\mathbb{D}(D)}(F) + \mathbb{W}_{\mathbb{D}(D)}(G) + \mathbb{W}_{\mathbb{D}(D)}(H) \\ &> 2 + \mathbb{W}_{\mathbb{D}(E)}(F) + \mathbb{W}_{\mathbb{D}(E)}(G) + \mathbb{W}_{\mathbb{D}(E)}(H) = \mathbb{W}_{\mathbb{D}(E)}(E) = \mathbb{W}(E). \end{aligned}$$

²This terminology is derived from linear logic, where proofs obtained by the promotion rule are usually called boxes

- Suppose that $D = \text{cut}_!(F, x.b_!(x, y.G)) \implies \text{cut}(F_\downarrow, y.\text{cut}_\#(F, x.G_\downarrow)) = E$. Then,

$$\begin{aligned}
\mathbb{D}(D) &= \max\{\mathbb{D}(F_\downarrow), \mathbb{D}(G_\downarrow)\} = \max\{\mathbb{D}(F), \mathbb{D}(F), \mathbb{D}(G)\} = \mathbb{D}(E); \\
\mathbb{W}(D) &= \mathbb{W}_{\mathbb{D}(D)}(D) = \text{FO}(x, b_!(x, y.G)) \cdot \mathbb{W}_{\mathbb{D}(D)}(F_\downarrow) + \mathbb{W}_{\mathbb{D}(D)}(b_!(x, y.G)) \\
&= \mathbb{W}_{\mathbb{D}(D)}(F) + \mathbb{W}_{\mathbb{D}(D)}(b_!(x, y.G)) = \mathbb{W}_{\mathbb{D}(D)}(F) + 1 + \mathbb{W}_{\mathbb{D}(D)}(G) \\
&\geq \mathbb{W}_{\mathbb{D}(E)}(F) + 1 + \mathbb{W}_{\mathbb{D}(E)}(G) \\
&> \mathbb{W}_{\mathbb{D}(E)}(F) + \mathbb{W}_{\mathbb{D}(E)}(G) = \mathbb{W}_{\mathbb{D}(E)}(F) + 0 \cdot \mathbb{W}_{\mathbb{D}(E)}(F) + \mathbb{W}_{\mathbb{D}(E)}(G) \\
&= \mathbb{W}_{\mathbb{D}(E)}(F) + \text{FO}(x, G) \cdot \mathbb{W}_{\mathbb{D}(E)}(F) + \mathbb{W}_{\mathbb{D}(E)}(G) \\
&= \mathbb{W}_{\mathbb{D}(E)}(E) = \mathbb{W}(E).
\end{aligned}$$

- Suppose that

$$D = \text{cut}_\#(F, x.b_\#(x, y.G)) \implies \text{cut}(F_\downarrow, y.\text{cut}_\#(F, x.G)) = E.$$

Then we can proceed exactly as in the previous case.

This concludes the proof. \square

Shift reduction, on the other hand, is *not* guaranteed to induce a strict decrease on the underlying weight which, however, cannot increase:

Proposition 3 *If $\Gamma; \Delta; \Theta \vdash D :: T$ and $D \mapsto E$, then $\Gamma; \Delta; \Theta \vdash E :: T$, $\mathbb{D}(E) \leq \mathbb{D}(D)$ and $\mathbb{W}(E) \leq \mathbb{W}(D)$.*

Proof. By induction on the proof that $D \mapsto E$. Some interesting cases:

- Suppose that

$$D = \text{cut}(!R(x_1, \dots, x_n, F), x.!L_!(x.G)) \mapsto !L_!(x_1.!L_!(x_2.\dots!L_!(x_n.\text{cut}_!(F, y.G))) = E.$$

Then,

$$\begin{aligned}
\mathbb{D}(D) &= \max\{\mathbb{D}(F), \mathbb{D}(G)\} = \mathbb{D}(E) \\
\mathbb{W}(D) &= \mathbb{W}_{\mathbb{D}(D)}(D) = \mathbb{D}(D) \cdot \mathbb{W}_{\mathbb{D}(D)}(F) + \mathbb{W}_{\mathbb{D}(D)}(G) \geq \text{FO}(y, G) \cdot \mathbb{W}_{\mathbb{D}(D)}(F) + \mathbb{W}_{\mathbb{D}(D)}(G) \\
&= \text{FO}(y, G) \cdot \mathbb{W}_{\mathbb{D}(E)}(F) + \mathbb{W}_{\mathbb{D}(E)}(G) = \mathbb{W}_{\mathbb{D}(E)}(E) = \mathbb{W}(E).
\end{aligned}$$

- Suppose that

$$D = \text{cut}(!R(x_1, \dots, x_n, F), x.!L_\#(x.G)) \mapsto !L_\#(x_1.!L_\#(x_2.\dots!L_\#(x_n.\text{cut}_\#(F, y.G))) = E.$$

Then we can proceed as in the previous case.

This concludes the proof. \square

Finally, equivalence leaves the weight unchanged:

Proposition 4 *If $\Gamma; \Delta; \Theta \vdash D :: T$ and $D \equiv E$, then $\Gamma; \Delta; \Theta \vdash E :: T$, $\mathbb{D}(E) = \mathbb{D}(D)$ and $\mathbb{W}(E) = \mathbb{W}(D)$.*

Proof. By induction on the proof that $D \equiv E$. Some interesting cases:

- Suppose that

$$D = \text{cut}(F, x.\text{cut}(G_x, y.H_y)) \equiv \text{cut}(\text{cut}(F, x.G_x), y.H_y) = E.$$

Then:

$$\begin{aligned}
\mathbb{D}(D) &= \max\{\mathbb{D}(F), \mathbb{D}(G_x), \mathbb{D}(H_y)\} = \mathbb{D}(E) \\
\mathbb{W}(D) &= \mathbb{W}_{\mathbb{D}(D)}(D) = \mathbb{W}_{\mathbb{D}(D)}(F) + \mathbb{W}_{\mathbb{D}(D)}(G_x) + \mathbb{W}_{\mathbb{D}(D)}(H_y) \\
&= \mathbb{W}_{\mathbb{D}(E)}(F) + \mathbb{W}_{\mathbb{D}(E)}(G_x) + \mathbb{W}_{\mathbb{D}(E)}(H_y) = \mathbb{W}_{\mathbb{D}(E)}(E) = \mathbb{W}(E).
\end{aligned}$$

- Suppose that

$$D = \text{cut}(F, x.\text{cut}(G, y.H_{xy})) \equiv \text{cut}(G, x.\text{cut}(F, y.H_{xy})) = E.$$

Then we can proceed as in the previous case.

- Suppose that

$$D = \text{cut}(F, x.\text{cut}_!(G, y.H_{xy})) \equiv \text{cut}_!(G, y.\text{cut}(F, x.H_{xy})) = E.$$

Then, since $\mathbb{FO}(y, F) = 0$,

$$\begin{aligned} \mathbb{D}(D) &= \max\{\mathbb{D}(F), \mathbb{D}(G), \mathbb{D}(H_{xy})\} = \mathbb{D}(E) \\ \mathbb{W}(D) &= \mathbb{W}_{\mathbb{D}(D)}(D) = \mathbb{W}_{\mathbb{D}(D)}(F) + \mathbb{FO}(y, H_{xy}) \cdot \mathbb{W}_{\mathbb{D}(D)}(G) + \mathbb{W}_{\mathbb{D}(D)}(H_{xy}) \\ &= \mathbb{W}_{\mathbb{D}(D)}(F) + \mathbb{FO}(y, \text{cut}(F, x.H_{xy})) \cdot \mathbb{W}_{\mathbb{D}(D)}(G) + \mathbb{W}_{\mathbb{D}(D)}(H_{xy}) \\ &= \mathbb{W}_{\mathbb{D}(E)}(F) + \mathbb{FO}(y, \text{cut}(F, x.H_{xy})) \cdot \mathbb{W}_{\mathbb{D}(E)}(G) + \mathbb{W}_{\mathbb{D}(E)}(H_{xy}) \\ &= \mathbb{W}_{\mathbb{D}(E)}(E) = \mathbb{W}(E). \end{aligned}$$

- Suppose that

$$D = \text{cut}_\#(F, x.\text{cut}(G_x, y.H_{xy})) \equiv \text{cut}(\text{cut}_\#(F, x.G_x), y.\text{cut}_\#(F, x.H_{xy})) = E.$$

Then,

$$\begin{aligned} \mathbb{D}(D) &= \max\{\mathbb{D}(F), \mathbb{D}(G_x), \mathbb{D}(H_{xy})\} = \mathbb{D}(E) \\ \mathbb{W}(D) &= \mathbb{FO}(x, \text{cut}(G_x, y.H_{xy})) \cdot \mathbb{W}_{\mathbb{D}(D)}(F) + \mathbb{W}_{\mathbb{D}(D)}(G_x) + \mathbb{W}_{\mathbb{D}(D)}(H_{xy}) \\ &= (\mathbb{FO}(x, G_x) + \mathbb{FO}(x, H_{xy})) \cdot \mathbb{W}_{\mathbb{D}(D)}(F) + \mathbb{W}_{\mathbb{D}(D)}(G_x) + \mathbb{W}_{\mathbb{D}(D)}(H_{xy}) \\ &= (\mathbb{FO}(x, G_x) \cdot \mathbb{W}_{\mathbb{D}(D)}(F) + \mathbb{FO}(x, H_{xy})) \cdot \mathbb{W}_{\mathbb{D}(D)}(F) + \mathbb{W}_{\mathbb{D}(D)}(G_x) + \mathbb{W}_{\mathbb{D}(D)}(H_{xy}) \\ &= \mathbb{W}_{\mathbb{D}(D)}(\text{cut}_\#(F, x.G_x)) + \mathbb{W}_{\mathbb{D}(D)}(\text{cut}_\#(F, x.H_{xy})) \\ &= \mathbb{W}_{\mathbb{D}(D)}(E) = \mathbb{W}_{\mathbb{D}(E)}(E) = \mathbb{W}(E). \end{aligned}$$

This concludes the proof. \square

Now, consider again the subject reduction theorem (Theorem 1): what it guarantees is that whenever $P \rightarrow Q$ and $\widehat{D} = P$, there is E with $\widehat{E} = Q$ and $D \leftrightarrow \implies \leftarrow E$. In view of the three propositions we have just stated and proved, it's clear that $\mathbb{W}(D) > E$. Altogether, this implies that $\mathbb{W}(D)$ is an upper bound on the number or internal reduction steps \widehat{D} can perform. But is $\mathbb{W}(D)$ itself bounded?

5.3 Bounding the Weight

What kind of bounds can we expect to prove for $\mathbb{W}(D)$? More specifically, how related are $\mathbb{W}(D)$ and $|\widehat{D}|$?

Lemma 5 *Suppose $\Gamma; \Delta; \Theta \vdash D :: T$. Then $\mathbb{D}(D) \leq |D|$.*

Proof. An easy induction on the structure of a type derivation π for $\Gamma; \Delta; \Theta \vdash D :: T$. \square

Lemma 6 *If $\Gamma; \Delta; \Theta \vdash D :: T$, then for every $n \geq \mathbb{D}(D)$, $\mathbb{W}_n(D) \leq |\widehat{D}| \cdot n^{\mathbb{B}(\widehat{D})+1}$.*

Proof. By induction on the structure of D . Some interesting cases:

- If $D = \text{cut}_!(D, x.E)$, then:

$$\begin{aligned}
\mathbb{W}_n(\text{cut}_!(D, x.E)) &= \mathbb{F}\mathbb{O}(x, E) \cdot (\mathbb{W}_n(D) + 1) + \mathbb{W}_n(E) \\
&\leq \mathbb{F}\mathbb{O}(x, E) \cdot (|D| \cdot n^{\mathbb{B}(D)+1} + 1) + |E| \cdot n^{\mathbb{B}(E)+1} \\
&\leq n \cdot |D| \cdot n^{\mathbb{B}(D)+1} + n + |E| \cdot n^{\mathbb{B}(E)+1} \\
&\leq |D| \cdot n^{\mathbb{B}(D)+2} + n^{\mathbb{B}(E)+1} + |E| \cdot n^{\mathbb{B}(E)+1} \\
&\leq (|D| + |E| + 1) \cdot n^{\max\{\mathbb{B}(D)+2, \mathbb{B}(E)+1\}} \\
&= |\text{cut}_!(D, x.E)| \cdot n^{\mathbb{B}(\text{cut}_!(D, x.E))}.
\end{aligned}$$

- If $D = !R(x_1, \dots, x_n, E)$, then:

$$\begin{aligned}
\mathbb{W}_n(!R(x_1, \dots, x_n, E)) &= n \cdot (\mathbb{W}_n(E) + 1) \\
&\leq n \cdot |E| \cdot n^{\mathbb{B}(E)+1} + n \\
&\leq |E| \cdot n^{\mathbb{B}(E)+2} + n^{\mathbb{B}(E)+2} \\
&= (1 + |E|) \cdot n^{\mathbb{B}(!R(x_1, \dots, x_n, E))+1} \\
&= !R(x_1, \dots, x_n, E) \cdot n^{\mathbb{B}(!R(x_1, \dots, x_n, E))+1}.
\end{aligned}$$

This concludes the proof. \square

5.4 Putting Everything Together

We now have almost all the necessary ingredients to obtain a proof of Proposition 1: the only missing pieces are the bounds on the size of any reducts, since the polynomial bounds on the length of internal reductions are exactly the ones from Lemma 6. Observe, however, that the latter induces the former:

Lemma 7 *Suppose that $P \rightarrow^n Q$. Then $|Q| \leq n \cdot |P|$.*

Proof. By induction on n , enriching the statement as follows: whenever $P \rightarrow^n Q$, both $|Q| \leq n \cdot |P|$ and $|R| \leq |P|$ for every subprocess R of Q in the form $!x(y).S$. \square

Let us now consider Theorem 2: how can we deduce it from Proposition 1? Everything boils down to show that for normal processes, the box-depth can be read off from their type. In the following lemma, $\mathbb{B}(A)$ and $\mathbb{B}(\Gamma)$ are the nesting depths of $!$ inside the type A and inside the types appearing in Γ (for every type A and context Γ).

Lemma 8 *Suppose that $\Gamma; \Delta; \Theta \vdash D :: x : A$ and that D is normal. Then $\mathbb{B}(\widehat{D}) = \max\{\mathbb{B}(\Gamma), \mathbb{B}(\Delta), \mathbb{B}(\Theta), \mathbb{B}(A)\}$.*

Proof. An easy induction on D . \square

6 Conclusions

In this paper, we introduced a variation on Caires and Pfenning's π DILL, called π DSL, being inspired by Lafont's soft linear logic. The key feature of π DSL is the fact that the amount of interaction induced

by allowing two processes to interact with each other is bounded by a polynomial whose degree can be “read off” from the type of the session channel through which they communicate.

What we consider the main achievement of this paper is definitely *not* the proof of these polynomial bounds, which can be obtained by adapting the ones in [6] or in [5], although this anyway presents some technical difficulties due to the low-level nature of the π -calculus compared to the lambda calculus or to higher-order π -calculus. Instead, what we found very interesting is that the operational properties induced by typability in π DSL, bounded interaction *in primis*, are not only very interesting and useful in practice, but different from the ones obtained in soft lambda calculi: in the latter, it’s the *normalization* time which is bounded, while here it’s the *interaction* time. Another aspect that we find interesting is the following: it seems that the constraints on processes induced by the adoption of the more stringent typing discipline π DSL, as opposed to π DILL, are quite natural and do not rule out too many interesting examples. In particular, the way sessions can be defined remains essentially untouched: what changes is the way sessions can be offered, i.e. the discipline governing the offering of multiple sessions by servers. All the examples in [1] and the one from Section 2 are indeed typable in π DSL.

Topics for future work include the accommodation of recursive types into π DSL. This could be easier than expected, due to the robustness of light logics to the presence of recursive types [3].

References

- [1] Luís Caires & Frank Pfenning (2010): *Session Types as Intuitionistic Linear Propositions*. In: *CONCUR 2010, LNCS 6269*. Springer, pp. 222–236.
- [2] Luís Caires, Bernardo Toninho & Frank Pfenning (2011): *Dependent Session Types via Intuitionistic Linear Type Theory*. In: *PPDP 2011*. ACM Press. To appear.
- [3] Ugo Dal Lago & Patrick Baillot (2006): *On light logics, uniform encodings and polynomial time*. *Mathematical Structures in Computer Science* 16(4), pp. 713–733.
- [4] Ugo Dal Lago & Paolo Di Giamberardino: *Soft Session Types (Long Version)*. Available at <http://arxiv.org/abs/1107.4478>.
- [5] Ugo Dal Lago, Simone Martini & Davide Sangiorgi (2010): *Light Logics and Higher-Order Processes*. In: *EXPRESS’10, EPTCS 41*. pp. 46–60.
- [6] Ugo Dal Lago, Andrea Masini & Margherita Zorzi (2010): *Quantum implicit computational complexity*. *Theor. Comput. Sci.* 411(2), pp. 377–409.
- [7] Jean-Yves Girard (1987): *Linear Logic*. *Theor. Comput. Sci.* 50, pp. 1–102.
- [8] Kohei Honda, Vasco Thudichum Vasconcelos & Makoto Kubo (1998): *Language Primitives and Type Discipline for Structured Communication-Based Programming*. In: *ESOP, LNCS 1381*. pp. 122–138.
- [9] Kohei Honda, Nobuko Yoshida & Marco Carbone (2008): *Multiparty asynchronous session types*. In: *POPL 2008*. ACM Press, pp. 273–284.
- [10] Yves Lafont (2004): *Soft linear logic and polynomial time*. *Theor. Comput. Sci.* 318(1-2), pp. 163–180.
- [11] Harry G. Mairson (1992): *A Simple Proof of a Theorem of Statman*. *Theor. Comput. Sci.* 103(2), pp. 387–394.
- [12] Dimitris Mostrous & Nobuko Yoshida (2007): *Two Session Typing Systems for Higher-Order Mobile Processes*. In: *TLCA 2007, LNCS 4583*. pp. 321–335.
- [13] Richard Statman (1979): *The Typed lambda-Calculus is not Elementary Recursive*. *Theor. Comput. Sci.* 9, pp. 73–81.