Full Abstraction for Linda

Cinzia Di Giusto and Maurizio Gabbrielli

Dip. Scienze dell'Informazione, Università di Bologna
Bologna, Italy
{digiusto, gabbri}@cs.unibo.it

Abstract. This paper investigates full abstraction of a trace semantics for two Linda-like languages. The first language provides primitives for adding and removing messages from a shared memory, local choice, parallel composition and recursion. The second one adds the possibility of checking for the absence of a message in the store. After having defined a denotational semantics based on traces, we obtain fully abstract semantics for both languages by using suitable abstractions in order to identify different traces which do not correspond to different operational behaviours.

1 Introduction

One of the fundamental purposes of a semantics is to provide a rigorous mean for proving the correctness of programs w.r.t. some behavioural specification. Several different tools (operational, denotational, algebraic and logic) can be used to this aim and ideally one would like to have a compositional and fully abstract semantics.

Compositionality is of course an important feature since it is the foundation for managing large systems complexity when considering program verification, analysis and (modular) design. Most of the above mentioned tools indeed allow to obtain rather easily a compositional semantics.

Full abstraction is also a desirable feature since it allows to simplify and “economize” as much as possible a semantics while preserving its correctness. However, in general this is a rather difficult target to achieve. To be more precise and to set the ground for the content of this paper, following [4,9,12] we can summarize the terms of the problem as follows. Given a language L , define a semantics that associates to each process (or program) P in L a set of observable properties $\mathcal{O}(P)$. This is usually done in operational terms by using a transition system and a suitable definition of $\mathcal{O}(P)$ which identifies computational aspects relevant for a specific class of applications. In case such semantics is compositional, i.e. if we can reconstruct $\mathcal{O}(P \text{ op } Q)$ from $\mathcal{O}(P)$ and $\mathcal{O}(Q)$ for any operator op of the language L , we have a satisfactory semantics, since the observational equivalence on processes induced by $\mathcal{O}(P)$ is preserved by contexts. More precisely, we have that $\mathcal{O}(P) = \mathcal{O}(Q)$ iff, for any context $C[\bullet]$, $\mathcal{O}(C[P]) = \mathcal{O}(C[Q])$.

However often this is not the case and in order to obtain a compositional semantics some richer semantic structures than those used in $\mathcal{O}(P)$ need to be considered. For example, as we will see in Section 4, typically pairs representing the input/output behaviour of a process are not sufficient to obtain compositionality and one has to use

traces. It can happen that these richer semantic structures “add too much” in the sense that the semantics $\llbracket \cdot \rrbracket$ based on them allows to distinguish processes which have the same behaviour w.r.t. $\mathcal{O}(P)$, under any possible context. In this case suitable abstractions must be used in $\llbracket \cdot \rrbracket$ in order to obtain a fully abstract result which, in general, can be stated as follows: $\llbracket P \rrbracket = \llbracket Q \rrbracket$ iff, for any context $C[\bullet]$, $\mathcal{O}(C[P]) = \mathcal{O}(C[Q])$ holds.

In this paper we investigate the full abstraction problem, as described above, for two variants of Linda. Linda is a programming paradigm [11] which allows interprocess communication through a shared data space, also called tuple space, where processes can post and retrieve messages (also called tuples). The shared memory paradigm offers some advantages since it decouples communication between processes: communication is in fact asynchronous and processes do not need to be aware of each other identity or location. Indeed, the Linda paradigm has received also a commercial interest, mainly due to the applications which use the Java Spaces from Sun Microsystems [10] and TSpaces from IBM [13] models, both based on Linda (a more detailed comparison of Linda implementations can be found in [19]). Distributed Linda-like languages have also been investigated. Notably, Klaim [17] is an implemented language based on the Linda paradigm where the central store is replaced by several distributed local stores and processes mobility among different locations is supported.

Fully abstract semantics based on traces for input/output observables have been studied many years ago for several concurrent languages, as we shall discuss in Section 6. However, to the best of our knowledge no one has yet addressed this problem for a Linda-like language.

Many different formalizations and variants of Linda have been defined. Here we use essentially the process-algebraic formalization of Linda introduced in [6,7] and we consider the very basic Linda dialects. The first one, which we call Linda-core, apart from the usual operators in process algebra (choice, parallel composition, recursion) contains the two Linda primitives *in* and *out* which allow to remove and add messages to the store, respectively. For Linda-core we define a compositional, fixpoint trace semantics which is correct but not fully abstract when considering the input/output pairs. Hence we introduce a suitable abstraction on traces and show that this allows us to obtain a fully abstract semantics. The second dialect (Linda-inp) enriches the syntax of Linda-core by allowing also a construct (*inp*) which allows to check the absence of information in the store. We prove that in this case a much simpler abstraction on traces is sufficient to obtain a full abstraction result. This accounts for the augmented expressive power of the language with *inp*, which can be formally proven by using the techniques in [6,20]. Unfortunately, due to the saturation operator, the fully abstract semantics are not compositional. This is unavoidable in our trace model, since the properties that we need to abstract depend on sets of traces (rather than on single ones). Of course this does not mean that in general a compositional fully abstract semantics based on traces does not exist. However, in case it existed, it would use traces substantially more complicated than ours.

The remainder of the paper is organized as follows. Section 2 introduces the Linda languages under consideration while Section 3 defines their denotational semantics. We then provide the fully abstract semantics for the core language in Section 4. Section 5

contains the main theorem on the full abstraction for the language extended with the *inp* primitive. Finally, Section 6 concludes by discussing some related works.

2 Preliminaries

In this section, following the process algebraic view of Linda proposed in [6] we recall the syntax of the Linda languages that we consider and their operational semantics.

2.1 Linda-Core

As previously mentioned, Linda is a paradigm which provides a simple model to describe communication between processes. The central notion in Linda is the one of *tuple space*. A tuple space is a shared data space (i.e. a common store) where all the *tuples* representing the information to be exchanged are stored. Here we shall abstract from the specific nature of tuples assuming that these are elementary messages. Communication is represented by the concurrent and asynchronous activity of several processes which add or remove messages from the common store. I.e. the sender dispatches a message through a non-blocking operation which adds the tuple in the tuple space. Then the message has an independent existence until a receiver retrieves and removes it from the shared space. Such kind of communication is called *generative* (see [11]).

Processes of the language Linda-core, denoted by $P, Q, \dots$, are then given by the following grammar:

$$P ::= 0 \mid out(a).P \mid in(a).P \mid P \mid P \mid P + P \mid recX.P \quad (1)$$

where we assume that $a \in Msg$ and Msg denotes the set of all possible messages (or tuples), ranged over by $a, b, \dots$.

Intuitively 0 represents the process that does nothing. Then the process $out(a).P$ adds the message a to the store and then behaves as P . The message a which has been added to the store will be visible to other processes only after the completion of the $out(a)$ action, however note that other interpretations are possible for this primitive (see [5]). If a is present in the tuple space, $in(a).P$ removes the message and then behaves as P . Otherwise if a is not present, the process $in(a).P$ is suspended until a becomes available in the store. The parallel construct $P \mid Q$ is interpreted in terms of interleaving. The process $P + Q$ can non-deterministically choose to behave either as P or as Q (hence we have a form of local choice). Finally we have the recursion operator where we assume that guarded recursion is used (i.e. the process $recX.X$ is not allowed).

The operational semantics of Linda-core is described by means of a transition system $T = (Conf, \rightarrow)$. Configurations $Conf$ are pairs of the form $\langle P, \mathcal{M} \rangle$ where P is a process and $\mathcal{M}$ is a multiset containing tuples, also called tuple space or store. The transition relation $\rightarrow \subseteq Conf \times Conf$ is the least relation satisfying the rules in Table 1, which should be self-explaining, provided we introduce the following notation.

Notation 1. To describe updates in the store we will use $\oplus$ and $\ominus$ to denote multisets union and difference, respectively. So $\mathcal{M} \oplus \{a\}$ means that a message (a tuple) ' a ' has been added to the store while $\mathcal{M} \ominus \{a\}$ indicates that a copy of ' a ' has been removed.

Table 1. An operational semantics An operational semantics for Linda-core

R1	$\langle out(a).P, \mathcal{M} \rangle \rightarrow \langle P, \mathcal{M} \oplus \{a\} \rangle$
R2	$\langle in(a).P, \mathcal{M} \oplus \{a\} \rangle \rightarrow \langle P, \mathcal{M} \rangle$
R3	$\frac{\langle P, \mathcal{M} \rangle \rightarrow \langle P', \mathcal{M}' \rangle}{\langle P \mid Q, \mathcal{M} \rangle \rightarrow \langle P' \mid Q, \mathcal{M}' \rangle} \text{ and } \frac{\langle Q, \mathcal{M} \rangle \rightarrow \langle Q', \mathcal{M}' \rangle}{\langle P \mid Q, \mathcal{M} \rangle \rightarrow \langle P \mid Q', \mathcal{M}' \rangle}$
R4	$\langle P + Q, \mathcal{M} \rangle \rightarrow \langle P, \mathcal{M} \rangle \text{ and } \langle P + Q, \mathcal{M} \rangle \rightarrow \langle Q, \mathcal{M} \rangle$
R5	$\frac{\langle P[recX.P/X], \mathcal{M} \rangle \rightarrow \langle P', \mathcal{M}' \rangle}{\langle recX.P, \mathcal{M} \rangle \rightarrow \langle P', \mathcal{M}' \rangle}$

Table 2. The rule for inp

R6	$\begin{array}{l} \langle inp(a)?P : Q, \mathcal{M} \oplus \{a\} \rangle \rightarrow \langle P, \mathcal{M} \rangle \\ \langle inp(a)?P : Q, \mathcal{M} \rangle \rightarrow \langle Q, \mathcal{M} \rangle \text{ provided } a \notin \mathcal{M} \end{array}$
----	---

A transition $\langle P, \mathcal{M} \rangle \rightarrow \langle Q, \mathcal{M}' \rangle$ then means that the process P reduces to Q , possibly by producing some changes in the store which evolves from $\mathcal{M}$ to $\mathcal{M}'$. A sequence of configurations is called run or computation. The reflexive transitive closure of $\rightarrow$ is denoted by $\Rightarrow$. By using the transition system described above we can characterize several different notions of observables. The ones we are interested in here consider simply the input/output behaviour of a process in terms of the tuple space. The input is therefore the initial tuple space, while the output is the final store produced by a process which cannot further proceed in the computation (denoted by $\nrightarrow$) either because it is suspended on an *in* operation or because it has consumed all the actions. More precisely we define the observables as follows.

Definition 1 (Observables $\mathcal{O}(P)$). Let P be a Linda process. We define:
 $\mathcal{O}(P) = \{(\mathcal{M}_1, \mathcal{M}_n) \mid \langle P, \mathcal{M}_1 \rangle \Rightarrow \langle P_n, \mathcal{M}_n \rangle \nrightarrow\}$

2.2 Linda-inp

We will now introduce a slightly different variant of Linda, called Linda-inp, obtained by adding a new operator $inp(a)?P : Q$ which allows also to check whether a message is not present in the store. More precisely, the previous construct checks whether the store contains the message a : if the message is present in the store then the process continues with P , otherwise with Q .

Therefore we will add to the grammar in (1) the following primitive:

$$P ::= inp(a)?P : P \quad (2)$$

The operational semantics for Linda-inp is obtained by (a transition system defined by) adding to the rules of Table 1 the rules contained in Table 2. The observables can be defined as before.

3 Denotational Semantics

It is easy to see that the operational semantics which associates to a process P its observables $\mathcal{O}(P)$ is not compositional. For example consider the processes $Q = out(b)$ and $P = out(a).in(a).out(b)$. Then $\mathcal{O}(P) = \mathcal{O}(Q)$ holds, however, considering the process $R = in(a).out(ok)$ we have that $(\emptyset, \{ok\}) \in \mathcal{O}(P \mid R) \setminus \mathcal{O}(Q \mid R)$ which means that the observables of a parallel composition cannot be obtained from the observables of the two processes being composed (in parallel). This problem is in general well known, in fact in order to obtain a compositional model more informative structures than input/output pairs have been used. In particular, models based on traces (or sequences) have been used for many concurrent languages, starting from the early works on dataflow languages [16], imperative ones [4] and concurrent constraint programming [9].

In the following we will define a compositional semantics which correctly models the $\mathcal{O}(P)$ observables and which is based on traces. This semantics is similar to those used for timed Linda in [8] (and therefore to that one of [9]), even though the technical treatment is different. In fact in [8], where maximal parallelism was assumed, the denotational model used traces of pairs of tuple spaces, representing the input and the output at each step of the computation. Here, due to the interleaving semantics and to local choice, this kind of sequences is not sufficient to obtain a correct model. Essentially the problem is that we have to distinguish the processes $out(a) \mid in(a)$ and $out(a).in(a) + in(a).out(a)$ (because when starting with an empty store the second process can produce an empty store as a result) and this cannot be done by using simply input/output pairs. Hence, here we consider a denotational model which associates to a process a set of sequences of the form $\alpha_1, \dots, \alpha_n$ where each α_i is an element of the set $\mathcal{A} = \{in(a), out(a), \overline{in}(a), \overline{inp}(a) \mid a \in Msg\}$ (where Msg denotes all the possible messages, as previously mentioned). The first two kinds of actions in $\mathcal{A}$ are obvious as they represent the corresponding operations on the store, $\overline{in}(a)$ and $\overline{inp}(a)$ are used to express absence of information. We denote with S the set of all possible sequences defined in this way.

We introduce now two denotational semantics (one for each language we are considering) based on traces which are compositional by construction. Such semantics are the least functions $\llbracket \cdot \rrbracket : Processes \rightarrow 2^S$, which satisfy the equations in Table 3 for Linda-core and the equations in Table 3 plus that in Table 4 for Linda-inp. The order on functions here is the one induced by set inclusion on the co-domain. Well known fixpoint results allow to obtain the semantics as the least fixpoint of the operators defined implicitly by the equations in the Tables.

3.1 Denotational Semantics for Linda-Core

The equations should be self-explanatory apart from a few details. The denotation of the 0 process is the empty sequence, while the equations D2 and D3 show the expected behaviour for the basic primitives. Note that in equation D3 we have two cases: the first one corresponds to the case in which a is present in the store, thus the computation can proceed (with the sequence s) after the in action. On the other hand, the $\overline{in}(a)$ action represents the absence of a in the store, in which case the computation terminates (the

Table 3. A denotational semantics for Linda-core

$$\text{D1 } \llbracket 0 \rrbracket = \{\epsilon\}$$

$$\text{D2 } \llbracket \text{out}(a).P \rrbracket = \{\text{out}(a) \cdot s \mid s \in \llbracket P \rrbracket\}$$

$$\text{D3 } \llbracket \text{in}(a).P \rrbracket = \{\text{in}(a) \cdot s \mid s \in \llbracket P \rrbracket\} \cup \{\overline{\text{in}}(a)\}$$

$$\text{D4 } \llbracket P \mid Q \rrbracket = \llbracket P \rrbracket \tilde{\mid} \llbracket Q \rrbracket$$

where the operator $\tilde{\mid}$ is inductively defined as follow:

$$(x \cdot s) \tilde{\mid} y = y \tilde{\mid} (x \cdot s) = \{(x \cdot t) \mid t \in s \tilde{\mid} y\} \cup \{y \cdot x \cdot s\}$$

$$(x \cdot s) \tilde{\mid} (y \cdot t) = (y \cdot t) \tilde{\mid} (x \cdot s) = \\ \{(x \cdot u) \mid u \in s \tilde{\mid} (y \cdot t)\} \cup \{(y \cdot u) \mid u \in (x \cdot s) \tilde{\mid} t\}$$

with $x, y \in \mathcal{A}$ and $s, t, u \in \mathcal{S}$.

$$\text{D5 } \llbracket P + Q \rrbracket = \llbracket P \rrbracket \cup \llbracket Q \rrbracket$$

$$\text{D6 } \llbracket \text{rec}X.P \rrbracket = \llbracket P[\text{rec}X.P/X] \rrbracket$$

process is suspended). The parallel operator is interpreted in terms of interleaving as usual, while since the choice is local, it can be modeled by a simple set union. Recursion is treated in the usual way.

In order to show that the denotational semantics is correct w.r.t. our notion of observables we define the evaluation of a trace as follows ($\uparrow$ means undefined).

Definition 2. Given a trace $s \in \mathcal{S}$ and a store $\mathcal{M}$, the function $\text{eval}_1(s, \mathcal{M})$ is defined by the following cases:

$$\text{eval}_1(\epsilon, \mathcal{M}) = \mathcal{M}$$

$$\text{eval}_1(\text{out}(x) \cdot t, \mathcal{M}) = \text{eval}_1(t, \mathcal{M} \oplus \{x\})$$

$$\text{eval}_1(\text{in}(x) \cdot t, \mathcal{M}) = \begin{cases} \text{eval}_1(t, \mathcal{M} \ominus \{x\}) & \text{if } x \in \mathcal{M} \\ \uparrow & \text{otherwise} \end{cases}$$

$$\text{eval}_1(\overline{\text{in}}(x) \cdot t, \mathcal{M}) = \begin{cases} \mathcal{M} & \text{if } x \notin \mathcal{M} \text{ and } t = \epsilon \\ \uparrow & \text{otherwise} \end{cases}$$

The correctness is then stated by the following proposition which can be proved by using a fixpoint characterization of the semantics $\llbracket \cdot \rrbracket$. This can be obtained by first considering an interpretation as a mapping $I : \text{Processes} \rightarrow 2^{\mathcal{S}}$ which associates to each process a denotation (i.e. a set of sequences). The set $\mathcal{I}$ of all the interpretations is easily seen to be a cpo with the ordering induced by $\subseteq$. An operator $\mathcal{F} : \mathcal{I} \rightarrow \mathcal{I}$ is obtained by substituting $\llbracket \cdot \rrbracket$ for $\mathcal{F}(I)$ in equations D1-D5 and in the left hand side of equation D6, and by replacing $\llbracket \cdot \rrbracket$ for I in the right hand side of equation D6. The semantics $\llbracket \cdot \rrbracket$ is then the least fixpoint of $\mathcal{F}$, which can be obtained as the least upper bound of $\{\mathcal{F}^n(\perp) \mid n \geq 0\}$, where $\perp$ is the least interpretation, $\mathcal{F}^0(\perp) = \perp$ and

Table 4. The equations for Linda-inp

$D7 \quad \llbracket inp(a)?P : Q \rrbracket = \{in(a) \cdot s \mid s \in \llbracket P \rrbracket\} \cup \{\overline{inp}(a) \cdot s \mid s \in \llbracket Q \rrbracket\}$
--

$\mathcal{F}^n(\perp) = \mathcal{F}(\mathcal{F}^{n-1}(\perp))$. This allows us to prove the proposition by induction on the structure of processes and on induction on the powers $\mathcal{F}^n(\perp)$ of the operator.

Proposition 1 (Correctness). *Given a Linda-core process P , $\mathcal{O}(P) = \{(\mathcal{M}_0, eval_1(s, \mathcal{M}_0)) \mid s \in \llbracket P \rrbracket \text{ and } eval_1(s, \mathcal{M}_0) \neq \uparrow\}$ holds.*

3.2 Denotational semantics for Linda-inp

When considering the Linda-inp language the denotational semantics can be obtained from Table 3 by adding the equation in Table 4. This difference w.r.t. the case of Linda core is due to the presence of the *inp*, which is described by Equation D7: since when a is present both the *inp*(a) and the *in*(a) construct are modeled in the same way, when a is not present we have to distinguish the two cases (by using $\overline{in}(a)$ and $\overline{inp}(a)$) since it would not be correct to use the evaluation given in Definition 2 for the *in*(a).

In order to prove the correctness of the model introduced above we need to add to $eval_1$ the new cases obtaining the evaluation function $eval_2$:

Definition 3. *Given a trace $s \in \mathcal{S}$ and a store $\mathcal{M}$, the function $eval_2(s, \mathcal{M})$ is defined by the following cases:*

$$eval_2(\overline{inp}(x) \cdot t, \mathcal{M}) = \begin{cases} eval_2(t, \mathcal{M}) & \text{if } x \notin \mathcal{M} \\ \uparrow & \text{otherwise} \end{cases}$$

$$eval_2(\alpha(x) \cdot t, \mathcal{M}) = eval_1(\alpha(x) \cdot t, \mathcal{M}) \text{ for } \alpha \neq \overline{inp}$$

Using the same technique of Proposition 1 it can be easily proved the following theorem that states the correctness of the denotational model:

Proposition 2 (Correctness). *Given a Linda-inp process P , $\mathcal{O}(P) = \{(\mathcal{M}_0, eval_2(s, \mathcal{M}_0)) \mid s \in \llbracket P \rrbracket\}$ holds.*

4 Full Abstraction for Linda-Core

The aim of this section is to obtain a fully abstract semantics for the Linda-core language. The semantics introduced in the previous section represents a too fine description of the actions that affect the store, since it records all the possible changes while the observables capture only the initial and the final state. It is therefore immediate to find processes which have a different denotation, while having the same input/output behaviour under any possible context.

In order to obtain full abstraction we saturate the denotational semantics by adding all those traces which, intuitively, represent a computation whose input/output behaviour, in any possible context, can be simulated by a trace which is already in the semantics. The formal definition is as follows.

Definition 4 (Saturation). Let $T \subseteq S$ be a set of traces. We define the saturation of T as the minimal set $Sat(T)$ which satisfies the following rules:

- i) if $s \in T$ then $s \in Sat(T)$
- ii) if $s \cdot out(a) \cdot t \cdot in(a) \cdot v \in Sat(T)$ then $s \cdot t \cdot v \in Sat(T)$
- iii) if $s \cdot out(a) \cdot t \cdot in(a) \cdot v \in Sat(T)$ then $s \cdot out(a) \cdot t \cdot in(a) \cdot out(a) \cdot in(a) \cdot v \in Sat(T)$
- iv) $s \in Sat(T)$ iff $s \cdot in(a) \cdot out(a) \in Sat(T)$
- v) if $s \cdot out(a) \cdot t \in Sat(T)$ then $s \cdot t' \in Sat(T)$ where $t' \in \{out(a) \cdot \tilde{t}\}$
- vi) all the traces in T of the form $t \cdot \overline{in}(a) \cdot u$ with $u \neq \epsilon$ are removed;

According to the previous definition in $Sat(T)$ we add all the traces which (i) are derived (inductively) from the traces in T by performing the following operations: (ii) Removing complementary actions $out(a)$ and $in(a)$ which appear, in this order, in different places of the sequence; it is rather clear that this does not change the operational behaviour described by the original sequence. (iii) Adding a “stuttering step” represented by a sequence $out(a) \cdot in(a)$ of two complementary actions is also allowed, provided that both these actions occur before (in this order) in the sequence. Intuitively, if the $out(a)$ action does not appear before in the sequence we cannot add it, since the presence of a could trigger some new computation; moreover, since the multiplicity of a message is relevant, also in case the sequence contains $out(a)$ and not $in(a)$ we cannot add the sequence $out(a) \cdot in(a)$ because after the added $out(a)$ we would have one more a than in the original sequence, which, again, could trigger new computations. (iv) Stuttering steps of the form $in(a) \cdot out(a)$ can be safely added and removed only at the end of a sequence. (v) As stated in [5] an output prefix $out(a).P$ is observably equivalent to $out(a) \mid P$, note that from this rule follows that the core-language cannot observe the order of appearance of messages. (vi) Finally, $\overline{in}(a)$ represents a process suspended because the message a is not present in the store, hence it is not correct to assume that other actions could take place afterwards. Clearly this is not anymore true (apart from rule (vi)) in presence of a construct which allows to check for absence of information, as we will see in the next section.

The fully abstract semantics is obtained by applying the saturation defined above to the semantics $\llbracket \cdot \rrbracket$.

In order to prove the full abstraction result we proceed by steps. First we prove that the abstraction introduced by Sat is correct (under any context) w.r.t. $\mathcal{O}(P)$. This result is obtained by first showing that the construction of $Sat(\llbracket P \rrbracket)$ does not add any trace that does not respect the observables of P . This is the content of the following Proposition, whose proof is immediate

Proposition 3. Given a process P , $\mathcal{O}(P) = \{(\mathcal{M}_0, eval_1(s, \mathcal{M}_0)) \mid s \in Sat(\llbracket P \rrbracket)\}$.

Now we are ready to state that the abstract (saturated) semantics is correct under any context w.r.t. the chosen observation criteria. A context $C[\bullet]$ is defined as a process with a hole, that is, a process where a subprocess is left unspecified. $C[P]$ is then the process obtained from $C[\bullet]$ by replacing $\bullet$ for the process P .

Theorem 2 (Correctness for Linda-core). Given two Linda-core process A and B , if $Sat(\llbracket A \rrbracket) = Sat(\llbracket B \rrbracket)$ then, for every context $C[\bullet]$, $\mathcal{O}(C[A]) = \mathcal{O}(C[B])$ holds.

Proof. We will first prove $\mathcal{O}(C[A]) \subseteq \mathcal{O}(C[B])$. Let $(\mathcal{M}_0, \mathcal{M}_1) \in \mathcal{O}(C[A])$ then following from Proposition 1 there exists $s \in \llbracket C[A] \rrbracket$ such that $\mathcal{M}_1 = eval_1(s, \mathcal{M}_0)$. Since the denotational semantics we provide is compositional $s = c \circ t$ for some suitable $\circ$, where $c \in \llbracket C[\bullet] \rrbracket$ and $t \in \llbracket A \rrbracket$.

Since $\llbracket A \rrbracket \subset Sat(\llbracket A \rrbracket) = Sat(\llbracket B \rrbracket)$ then $t \in Sat(\llbracket B \rrbracket)$ therefore two cases could arise: (1) $t \in \llbracket B \rrbracket$ hence $s \in \llbracket C[B] \rrbracket$ and $(\mathcal{M}_0, \mathcal{M}_1) \in \mathcal{O}(C[B])$. (2) $t \notin \llbracket B \rrbracket$ then there exists $u \in \llbracket B \rrbracket$ such that u is derived from t following the rules in definition 4 and $eval_1(t, \mathcal{M}_0) = eval_1(u, \mathcal{M}_0)$. Hence by induction on the structure of c it can be easily proved that $eval_1(c \circ u, \mathcal{M}_0) = \mathcal{M}_1$ and therefore $(\mathcal{M}_0, \mathcal{M}_1) \in \mathcal{O}(C[B])$.

The other set inclusion $\mathcal{O}(C[B]) \subseteq \mathcal{O}(C[A])$ is symmetrical. $\square$

To obtain full abstraction we need now to prove the converse of the above theorem. This is the central result of this section and is the content of the following.

Theorem 3. *Given two Linda-core processes A and B , if $Sat(\llbracket A \rrbracket) \neq Sat(\llbracket B \rrbracket)$ then there exists a context $C[\bullet]$ such that $\mathcal{O}(C[A]) \neq \mathcal{O}(C[B])$.*

Proof. Suppose that there exists $t \in Sat(\llbracket A \rrbracket) \setminus Sat(\llbracket B \rrbracket)$ and consider a generic $s \in Sat(\llbracket B \rrbracket)$ (thus $t \neq s$). From the definition of Sat it follows that we can choose s and t as the shortest sequences such that: (i) they do not contain sub-sequences of the form $out(x) \cdot u \cdot in(x) \cdot out(x) \cdot in(x)$, (ii) they do not contain suffixes of the form $in(x) \cdot out(x)$, (iii) every output appears as soon as possible and (iv) between two consecutive inputs the outputs are ordered in lexicographic order.

Then assume that t and s have the following form: $t = r \cdot \alpha(x) \cdot t_1$, $s = r \cdot \beta(y) \cdot s_1$ where the common prefix r can also be empty and $\alpha, \beta \in \mathcal{A}$ with $\alpha \neq \beta$.

The proof is by cases, where we analyze the first couple of different actions α and β . In each case we will construct a context $C[\bullet]$ which allows to distinguish A and B (that is, a context such that $\mathcal{O}(C[A]) \neq \mathcal{O}(C[B])$). In the proof we will use the following notation: if $in(a_1), in(a_2), \dots, in(a_n)$ are all the input actions which appear, in this order, in the sequence r (which can also contain other out actions), then $InComp(r)$ denotes the sequence $out(a_1) \cdot out(a_2) \cdots out(a_n)$: intuitively this sequence is a sort of complement (w.r.t. in actions) of r which allows to proceed in the computation when composed in parallel with r . Furthermore, in order to further simplify the notation, in the following we will use these assumptions:

$$c_1 = InComp(r)$$

c_2 is a sequence consisting of as many $in(x)$ as the $out(x)$ in r

c_3 is a sequence consisting of as many $in(y)$ as the $out(y)$ in r

We have then the following cases:

1. let $\beta(y) \cdot s_1 = \epsilon$, thus $t = r \cdot \alpha(x) \cdot t_1$ and $s = r$. Depending on t we can construct the following distinguishing contexts $C[\bullet]$:
 - (a) if $t = r \cdot out(x) \cdot t_1$ then $C[\bullet] = \bullet \mid c_1.c_2.in(x).out(ok)$;
 - (b) if $t = r \cdot in(x) \cdot t_1$ noticing that $t_1 \neq out(x)$, the following context can be provided $C[\bullet] = \bullet \mid c_1.out(x).InComp(t_1)$.

The symmetric case is completely analogous.

2. $\alpha(x) = in(x)$ and $\beta(y) = in(y)$ ($x \neq y$) then in order to separate the two processes we need to make further distinctions (note that by construction $t_1 \neq out(x)$):
 - (a) if $eval_1(t_1, \emptyset) \neq \{x\}$ then $C[\bullet] = \bullet \mid c_1.c_3.out(x)$
 - (b) if $eval_1(t_1, \emptyset) = \{x\}$ and the actions $out(y)$, $in(y)$ do not appear in t_1 then $C[\bullet] = \bullet \mid c_1.c_3.out(x).InComp(t_2)$
 - (c) otherwise since $out(y)$ appears in t_1 , it can be provided the following context $C[\bullet] = \bullet \mid c_1.c_3.out(x).in(y).out(y).out(y)$.
3. $\alpha(x) = out(x)$ and $\beta(y) = in(y)$ or vice versa: then it can be easily shown that the context $C[\bullet] = \bullet \mid c_1.c_2.c_3.in(x).out(ok)$ allows to distinguish A and B .
4. $\alpha(x) = out(x)$ and $\beta(y) = out(y)$ (with $x \neq y$). By hypothesis we can choose $t = r \cdot out(x) \dots in(v) \dots$ and $s = r \cdot out(y) \dots in(w) \dots$ where $in(v)$ and $in(w)$ are the first input actions after $out(x)$ and $out(y)$ respectively. Moreover $out(x)$ does not appear before $in(w)$ in s . Then two cases could arise if $v \neq x$ then the context $C[\bullet] = \bullet \mid c_1.c_4.c_5$ where c_4 and c_5 are sequences of as many $in(v)$ and $in(w)$ as the $out(v)$ and $out(w)$ that precedes the two input actions respectively. Instead if $v = x$ then we can safely assume $in(w)$ does not appear in s and the context $C[\bullet] = \bullet \mid c_1.c_5.InComp(t_1)$ can distinguish the two processes.
5. There are some remaining cases, where the two sequences are different because of a $\overline{in}$ action. However, due to the construction of our semantics, $r \cdot \overline{in}(x) \in Sat_2(\llbracket A \rrbracket)$ iff $r \cdot in(x) \cdot s \in Sat_2(\llbracket A \rrbracket)$. Therefore we can omit to consider the sequence $r \cdot \overline{in}(x)$ and just consider the case of the sequence $r \cdot in(x) \cdot s$, which is included above.

This completes the proof. $\square$

The previous two theorems can be summarized in the following corollary.

Corollary 1 (Full Abstraction for Linda-core). *Given two Linda-core processes A and B , $Sat(\llbracket A \rrbracket) = Sat(\llbracket B \rrbracket)$ iff, for any context $C[\bullet]$, $\mathcal{O}(C[A]) = \mathcal{O}(C[B])$ holds.*

5 Full Abstraction for Linda-inp

Now we move to consider the language Linda-inp where we can test for the absence of a message in the store by using the primitive *inp*. As underlined in the introduction, such a possibility augments the expressive power of the language. In semantic terms this means that we can construct more powerful contexts, thus allowing to discriminate processes which were identified by Linda-core contexts. As a simple example, consider the two processes $A = out(a).out(b)$ and $B = out(b).out(a)$. These processes cannot be distinguished (w.r.t. the observables $\mathcal{O}$) by any Linda-core contexts, indeed the corresponding traces $out(a) \cdot out(b)$ and $out(b) \cdot out(a)$ are identified by the saturation operation. However, the context $C[\bullet] = \bullet \mid in(a).(inp(b)?out(nok) : out(ok))$ allows to distinguish them, since it allows to check that a is present and b is absent in the store. Indeed we have that $(\emptyset, ok \in \mathcal{O}(C[A]) \setminus \mathcal{O}(C[B]))$. This example shows that a fully abstract semantics for Linda-inp must induce a finer equivalence on processes than Sat or, in other terms, that a less abstract operation has to be used to saturate sequences. However the Denotational semantics provided in Section 3.2 is not fully abstract. In fact, consider the two processes $A = inp(a)?\mathbf{0} : \mathbf{0}$ and $B = in(a) + A$: these two processes cannot be distinguished by any context, yet they have a different denotational semantics. Thus we need the following definition.

Definition 5 (Saturation for Linda-inp). Let $T \subseteq \mathcal{S}$ be a set of traces. We define the *inp-saturation* of T as the set $Sat_2(T)$ which is obtained by performing the following steps (in this order) on T :

1. all the traces in T of the form $t \cdot \overline{in}(a) \cdot u$ with $u \neq \epsilon$ are removed;
2. all the $\overline{in}(x)$ actions in all traces are replaced by $\overline{inp}(x)$ (for any x).

Condition 1 ensures that we obtain correct traces once we have performed the transformation in 2. In fact, $\overline{in}(a)$ comes from the evaluation of $in(a)$, when a is not present. Since such an evaluation is suspended, it is not correct to assume that some action can be performed later. Thus, before transforming $\overline{in}(a)$ into $\overline{inp}(a)$ (and therefore moving from the $eval_1$ of Definition 2 to $eval_2$ of Definition 3) we have to delete these traces. The correctness of the saturation is stated by the following proposition which can be easily proved.

Proposition 4. Given a process P , $\mathcal{O}(P) = \{(\mathcal{M}_0, eval_2(s, \mathcal{M}_0)) \mid s \in Sat_2(\llbracket P \rrbracket)\}$

Theorem 4. Given two Linda-inp processes A and B , if $Sat_2(\llbracket A \rrbracket) = Sat_2(\llbracket B \rrbracket)$ then, for every context $C[\bullet]$, $\mathcal{O}(C[A]) = \mathcal{O}(C[B])$ holds.

Proof. We will first prove $\mathcal{O}(C[A]) \subseteq \mathcal{O}(C[B])$. Let $(\mathcal{M}_0, \mathcal{M}_1) \in \mathcal{O}(C[A])$ then following from Proposition 2 there exists $s \in \llbracket C[A] \rrbracket$ such that $\mathcal{M}_1 = eval_2(s, \mathcal{M}_0)$. Since the denotational semantics we provide is compositional $s = c \tilde{o} t$ for some suitable $\tilde{o}$, where $c \in \llbracket C[\bullet] \rrbracket$ and $t \in \llbracket A \rrbracket$.

Applying the rules in definition 5 we can construct a trace t' s.t. $eval_2(t, \mathcal{M}_0) = eval_2(t', \mathcal{M}_0)$. Hence $t' \in Sat_2(\llbracket A \rrbracket)$. Since $Sat_2(\llbracket A \rrbracket) = Sat_2(\llbracket B \rrbracket)$, $t' \in Sat_2(\llbracket B \rrbracket)$ two cases could arise: (1) $t' \in \llbracket B \rrbracket$ hence $s \in \llbracket C[B] \rrbracket$ and $(\mathcal{M}_0, \mathcal{M}_1) \in \mathcal{O}(C[B])$. Or (2) $t' \notin \llbracket B \rrbracket$ therefore there exists $u \in \llbracket B \rrbracket$ where some of the actions $\overline{in}$ have been replaced with $\overline{inp}$ and $eval_2(t, \mathcal{M}_0) = eval_2(u, \mathcal{M}_0)$. Hence by induction on the structure of c it can be easily proved that $eval_2(c \tilde{o} u, \mathcal{M}_0) = \mathcal{M}_1$ and therefore $(\mathcal{M}_0, \mathcal{M}_1) \in \mathcal{O}(C[B])$.

The other set inclusion $\mathcal{O}(C[B]) \subseteq \mathcal{O}(C[A])$ is symmetrical. $\square$

Theorem 5. Given two Linda-inp processes A and B , if $Sat_2(\llbracket A \rrbracket) \neq Sat_2(\llbracket B \rrbracket)$ then there exists a context $C[\bullet]$ such that $\mathcal{O}(C[A]) \neq \mathcal{O}(C[B])$.

Proof. Suppose that there exists $t \in Sat_2(\llbracket A \rrbracket) \setminus Sat_2(\llbracket B \rrbracket)$ and consider a generic $s \in Sat_2(\llbracket B \rrbracket)$. Since $s \neq t$ by hypothesis, let $t = r \cdot \alpha_1(x_1) \cdots \alpha_n(x_n)$ and $s = r \cdot \beta_1(y_1) \cdots \beta_m(y_m)$ where the common prefix r can also be empty and $\alpha_1, \dots, \alpha_n, \beta_1, \dots, \beta_m \in \mathcal{A}$ with $\alpha_1 \neq \beta_1$.

The proof is by cases, where we analyze the first different actions α_1 and β_1 in the sequences t and s . In each case we will construct a context $C[\bullet]$ which allows to distinguish A and B (that is, a context such that $\mathcal{O}(C[A]) \neq \mathcal{O}(C[B])$). As in the proof of Theorem 3, if $in(a_1), in(a_2), \dots, in(a_n)$ are all the input actions which appear, in this order, in the sequence r then $InComp(r)$ denotes the sequence $out(a_1) \cdot out(a_2) \cdots out(a_n)$. Furthermore, in order to further simplify the notation, in the following we will use these assumptions:

$$c_1 = InComp(r)$$

c_2 is a sequence consisting of as many $in(a)$ as the $out(a)$ in r

c_3 is a sequence consisting of as many $in(b)$ as the $out(b)$ in r

We have then the following cases:

1. $t = r \cdot out(a) \cdot t_1$ and $s = r$; In this case $C[\bullet] = \bullet \mid c_1.c_2.in(a).out(ok)$ allows to distinguish A and B .
2. $t = r \cdot in(a) \cdot t_1$ and $s = r$; then $C[\bullet] = out(a).\bullet \mid c_1.c_2.inp(a)?out(ok) : out(no)$ is the distinguishing context.
3. $t = r \cdot out(a) \cdot t_1$ and $s = r \cdot out(b) \cdot s_1$. We have the following sub-cases:
 - (a) If the number of $out(a)$ in t is different from the number of $out(a)$ in s then it can be easily proved that there is a context that distinguishes the two programs (essentially it is a context that *counts* the occurrences of the $out(a)$). Similarly if we are considering the b 's. The following is an example.

Example 1. If $t = out(a) \cdot in(a) \cdot out(a) \cdot out(b)$ and $s = out(b) \cdot out(a)$ then we can build the distinguishing context

$$C[\bullet] = \bullet \mid in(a).out(a).inp(a)?out(ok) : out(no)$$

- (b) Now suppose that the number of $out(a)$ (or $out(b)$) is the same in t and s . If in t_1 or in s_1 there is an input action again it is easy to provide a distinguishing context, either by blocking the execution of the rest of the trace after the input or by querying the store for the presence/absence of messages in the store. The following provide an example.

Example 2. If $t = out(a) \cdot in(b) \cdot out(b)$ and $s = out(b) \cdot in(b) \cdot out(a)$ then we can consider the distinguishing context

$$C[\bullet] = \bullet \mid in(a).out(b).inp(b)?out(ok) : out(no)$$

- (c) If in t_1 and in s_1 there are only outputs then either there is an output action which is not present in one of the two traces, and in this case it is straightforward to build a distinguishing context, or the output actions of a sequence are a permutation of output actions of the other sequence; also in this case it is easy to construct a context that distinguishes the two processes by checking the presence of a message and the absence of the other one, as shown by the following.

Example 3. If $t = out(a) \cdot out(b)$ and $s = out(b) \cdot out(a)$ then the distinguishing context $C[\bullet] = \bullet \mid in(a).inp(b)?out(ok) : out(no)$ (as seen in the initial part of this Section).

4. $t = r \cdot out(a) \cdot t_1$, $s = r \cdot in(b) \cdot s_1$ and $s' = r \cdot \overline{inp}(b) \in Sat_2(\llbracket B \rrbracket)$. It suffices to consider $C[\bullet] = \bullet \mid c_1.c_2.c_3.in(a).out(ok)$.
5. $t = r \cdot out(a) \cdot t_1$, $s = r \cdot in(b) \cdot s_1$ and $s' = r \cdot \overline{inp}(b) \cdot s_2 \in Sat_2(\llbracket B \rrbracket)$. The following situations may arise:

- (a) if $out(a) \notin s_2$ then $C[\bullet] = \bullet \mid c_1.c_2.c_3.in(a).out(ok)$;
- (b) if $in(b) \notin t_1$ then $C[\bullet] = out(b).\bullet \mid c_1.inp(b)?out(ok) : out(no)$;
- (c) otherwise the only significant case is when $s' = r \cdot \overline{inp}(b) \cdot out(a) \cdot t_1$ and therefore a suitable context can be constructed observing that the order of the actions is different (i.e. b is consumed in two different positions). This is shown in the following.

Example 4. If $t = out(a)$ and $s = in(b) \cdot out(b) \cdot out(a)$ recalling that $s' = \overline{inp}(b) \cdot out(a)$ we can build the distinguishing context $C[\bullet] = out(b).\bullet \mid in(b).out(b).inp(b)?out(ok) : out(no)$

- 6. $t = r \cdot in(a) \cdot t_1$, $s = r \cdot in(b) \cdot s_1$ and $s = r \cdot \overline{inp}(b) \in Sat_2(B)$. In this case $C[\bullet] = out(a).\bullet \mid c_1.c_2.c_3.inp(a)?out(ok) : out(no)$.
- 7. $t = r \cdot in(a) \cdot t_1$, $s = r \cdot in(b) \cdot s_1$ and $s' = r \cdot \overline{inp}(b) \cdot s_2 \in Sat_2(\llbracket B \rrbracket)$. We should here distinguish between the following further cases
 - (a) if $out(a) \notin t_1$ and $in(a) \notin s_2$ then $C[\bullet] = out(a).\bullet \mid c_1.c_3$;
 - (b) otherwise the worst possible scenario happens when $s_2 = in(a) \cdot t_1$ and t_1 and s_1 are “symmetrical” in a and b . As already shown in some preceding cases, when the order of the actions changes it is always possible to find a distinguishing context. This is shown in the following, last example.

Example 5. Given $A = inp(a)?(out(a).in(b).out(b)) : (in(b).out(b))$, and $B = inp(b)?(out(b).in(a).out(a)) : (in(a).out(a))$ thus $Sat_2(\llbracket A \rrbracket) = \{in(a) \cdot out(a) \cdot in(b) \cdot out(b), \overline{inp}(a) \cdot in(b) \cdot out(b), \dots\}$ and $Sat_2(\llbracket B \rrbracket) = \{in(b) \cdot out(b) \cdot in(a) \cdot out(a), \overline{inp}(b) \cdot in(a) \cdot out(a), \dots\}$ and the following context can distinguish between the two programs: $C[\bullet] = \bullet \mid inp(a)?out(ok1) : (inp(b)?out(ok2) : out(no))$.

- 8. There are some remaining cases, where the two sequences are different because of a $\overline{inp}$ action. However, due to the construction of our semantics, $r \cdot \overline{inp}(x) \in Sat_2(\llbracket A \rrbracket)$ iff $r \cdot in(x) \cdot s \in Sat_2(\llbracket A \rrbracket)$. Therefore we can omit to consider the sequence $r \cdot \overline{inp}(x)$ and just consider the case of the sequence $r \cdot in(x) \cdot s$, which is included above.

This completes the proof. □

The previous two theorems can be summarized in the following corollary.

Corollary 2 (Full Abstraction for Linda-inp). *Given two Linda-inp processes A and B , $Sat_2(\llbracket A \rrbracket) = Sat_2(\llbracket B \rrbracket)$ iff, for any context $C[\bullet]$, $\mathcal{O}(C[A]) = \mathcal{O}(C[B])$ holds.*

6 Conclusions and Related Work

We have studied the full abstraction problem for two variants of the Linda paradigm. For the first one, the core Linda language, we have provided a trace semantics which is fully abstract w.r.t. the input/output notion of observables. This has been obtained by using a suitable abstraction in order to identify different traces which do not represent

meaningful operational differences. The second language, Linda-*inp*, allows also checking for the absence of information. The augmented expressive power of this language permits us to obtain a full abstraction result by using a much simpler abstraction.

In the specific context of Linda, full abstraction has been previously investigated by [3] which used also techniques inspired by [12]. The results in [3] are completely different from ours, since in such a paper a semantics based on sequences is shown to be fully abstract with respect to a notion of observable which consider traces of computations. We prefer to consider a coarser notion of observables, consisting in the input/output behaviour, which accounts for a “black box” use of processes. Clearly our notion of observables leads to a more difficult full abstraction result, being the denotational model based on traces.

Results similar to ours have been obtained in the context of concurrent constraint programming (CCP) by De Boer and Palamidessi [9], however this language differs from Linda since it does not allow to remove information from the store. This monotonic nature of CCP makes its semantic treatment simpler, hence the results in [9] cannot be applied directly to the languages we consider here. Also Brookes [4] provides a trace model and a full abstraction result for a shared variable parallel language which is substantially different from Linda. The same applies to the results in [12].

More generally, full abstraction results have been provided for many concurrent languages and in quite various settings, which however are different from the case we consider here. In fact, even though our core Linda language can be seen as a fragment of asynchronous CCS (and therefore of asynchronous π -calculus), all the full abstraction results available for these languages consider different observational equivalences from ours. Probably the closer work in this sense is [2], where full abstraction of a trace semantics w.r.t. may testing equivalence has been studied. Note however that may testing is different from the observational equivalence that we consider (which is based on the input-output behaviour). For example, the processes $in(a).in(b)$ and $in(b).in(a)$ are may testing equivalent (see [2]) while they are not equivalent in our case, since they can be distinguished by the context $out(a)$.

Several other papers consider barbed equivalences and their relations with bisimulation, (notably [1] for asynchronous π -calculus and [6] for Linda-like process algebras) which, as previously mentioned, are completely different from the equivalence we consider. It is also worth noticing that the construct *inp*, which is not available in π -calculus and in CCS, change considerably the semantics of the language, thus for Linda-*inp* one cannot use existing results for CCS or π -calculus. For example, [6] shows that in presence of *inp* the coarse congruence contained in barbed equivalence is a new, specific congruence called *inp*-bisimulation while for the core language it is the usual one.

Recently, full abstraction results for π -calculus with contextual equivalence [18] and for Java-like languages with testing equivalence have been obtained in [15] (by considering weak bisimulation) and in [14] (by using a model based on traces). Also in these cases the considered equivalences are different from ours.

Our results can be extended along several lines. We have described a fully abstract semantics which is not compositional since the abstractions that we need on sequences are inherently non compositional. It would therefore be interesting to determine whether a simple, compositional, fully abstract semantics based on sequences actually exists.

Secondly we could investigate some other of the (many) dialects of Linda which exist in the literature. In particular we are planning to investigate the case of the language Klaim [17], which supports distribution and mobility. Finally it would be interesting to consider full abstraction results for other notions of observational equivalences.

References

1. Amadio, R.M., Castellani, I., Sangiorgi, D.: On bisimulations of the asynchronous π -calculus. *Theor. Comput. Sci.* 195(2), 291–324 (1998)
2. Boreale, M., Nicola, R.D., Pugliese, R.: Trace and testing equivalence on asynchronous processes. *Inf. Comput.* 172(2), 139–164 (2002)
3. Brogi, A., Jaquet, J.-M.: Modeling coordination via asynchronous communication. In: *Proceedings of the 2nd COORD 1997*, London, UK, pp. 238–255. Springer, Heidelberg (1997)
4. Brookes, S.D.: Full abstraction for a shared-variable parallel language. *Information and Computation* 127(2), 145–163 (1996)
5. Busi, N., Gorrieri, R., Zavattaro, G.: Three semantics of the output operation for generative communication. In: *Proceedings of the 2nd COORD 1997*, London, UK, pp. 205–219. Springer, Heidelberg (1997)
6. Busi, N., Gorrieri, R., Zavattaro, G.: A process algebraic view of linda coordination primitives. *Theor. Comput. Sci.* 192(2), 167–199 (1998)
7. Busi, N., Gorrieri, R., Zavattaro, G.: On the turing equivalence of linda coordination primitives. *Theor. Comput. Sci.* 230(1-2), 260–261 (2000)
8. de Boer, F.S., Gabbrielli, M., Meo, M.C.: A timed linda language and its denotational semantics. *Fundamenta Informaticae* 63(4) (2004)
9. de Boer, F.S., Palamidessi, C.: A fully abstract model for concurrent constraint programming. In: *Proceedings TAPSOFT/CAAP 1991*, vol. 1, pp. 296–319. Springer, New York (1991)
10. Freeman, E., Arnold, K., Hupfer, S.: *JavaSpaces Principles, Patterns, and Practice*. Addison-Wesley Longman Ltd, Essex, UK (1999)
11. Gelernter, D.: Generative communication in linda. *ACM Trans. Program. Lang. Syst.* 7(1), 80–112 (1985)
12. Horita, E., de Bakker, J.W., Rutten, J.J.M.M.: Fully abstract denotational models for nonuniform concurrent languages. *Inf. Comput.* 115(1), 125–178 (1994)
13. IBM. Tspaces, <http://www.almaden.ibm.com/cs/TSpaces/index.html>
14. Jeffrey, A., Rathke, J.: Java jr.: Fully abstract trace semantics for a core java language. In: Sagiv, M. (ed.) *ESOP 2005*. LNCS, vol. 3444, pp. 423–438. Springer, Heidelberg (2005)
15. Jeffrey, A.S.A., Rathke, J.: Full abstraction for polymorphic pi-calculus. *Theoretical Computer Science* (to appear, 2007)
16. Jonsson, B.: A model and proof system for asynchronous networks. In: *Proc. of the 4th ACM symposium on Principles of distributed computing*, pp. 49–58. ACM Press, New York (1985)
17. Nicola, R.D., Ferrari, G.L., Pugliese, R.: Klaim: A kernel language for agents interaction and mobility. *IEEE Transactions on Software Engineering* 24(5), 315–330 (1998)
18. Pierce, B.C., Sangiorgi, D.: Behavioral equivalence in the polymorphic pi-calculus. In: *Proceedings of the 24th ACM SIGPLAN-SIGACT*, pp. 242–255. ACM Press, New York (1997)
19. Wells, G., Clayton, P., Chalmers, A.G.: A Comparison of Linda Implementations in Java. In: *Communicating Process Architectures 2000*, september 2000, pp. 63–76 (2000)
20. Zavattaro, G.: Towards a hierarchy of negative test operators for generative communication. *Electr. Notes Theor. Comput. Sci.* 16(2) (1998)