# cJoin: Join with communicating transactions[†]

R O B E R T O   B R U N I[‡], H E R N Á N   M E L G R A T T I[§]

and U G O   M O N T A N A R I[‡]

[‡]*Dipartimento di Informatica, Università di Pisa*
*Largo Bruno Pontecorvo 3, I-56127 Pisa, Italy*
*Email:* {`bruni; ugo`}`@di.unipi.it`
[§]*Departamento de Computación, FCEyN, Universidad de Buenos Aires - CONICET*
*Pabellón I, Ciudad Universitaria, (C1428EGA) Buenos Aires, Argentina*
*Email:* `hmelgra@dc.uba.ar`

This paper proposes a formal approach to the design and programming of *long running transactions* (LRTs). We exploit techniques from process calculi to define cJoin, which is an extension of the Join calculus with few well-disciplined primitives for LRT. Transactions in cJoin are intended to describe the transactional interaction of several partners, under the assumption that any partner executing a transaction may communicate only with other transactional partners. In such case, the transactions run by any party are bound to achieve the same outcome (i.e., all succeed or all fail). Hence, a distinguishing feature of cJoin, called *dynamic joinability*, is that ongoing transactions can be merged to complete their tasks and when this happens either all succeed or all abort. Additionally, cJoin is based on compensations i.e., partial executions of transactions are recovered by executing user-defined programs instead of providing automatic rollback. The expressiveness and generality of cJoin is demonstrated by many examples addressing common programming patterns. The mathematical foundation is accompanied by a prototype language implementation, which is an extension of the JoCaml compiler.

## 1. Introduction

The ultimate goal of *service oriented architecture* (SOA) is to make it possible to develop new components and applications (now services) just by assembling existing ones. Many recent efforts, strongly pushed by large industrial consortia, have given birth to several (proposals for) programming/description languages tailored to the specification of web service integration, generally known as *web service composition languages* (WSCL), like XLANG (Thatte 2001), WSFL (Leymann 2001), WS-BPEL (BPEL 2003), WS-CDL (WSCDL 2004), WSCI (WSCI 2002) and BPMN (BPMN 2010). WSCLs address aggregation by following two complementary approaches:

— *Orchestration*: a composite service consists essentially of a unique program (usually known as *orchestrator*) that coordinates the execution of all components, while involved

services are neither aware of the fact they are taking part in a larger process nor of its structure and goal. The application logic of a composite service relies on the orchestrator, which is responsible for interacting with (i.e., invoking) all involved components in the right order. For this reason, orchestration is appropriate for specifying intra-organisation (or private) processes, whose application details may be completely known and whose execution may be coordinated in a centralised way. Typical orchestration languages are XLANG, WSFL, and (executable processes of) WS-BPEL and BPMN.

— *Choreography*: choreographies do not rely on a centralised coordinator since they are intended to facilitate the integration of business processes that spawn over different organisations. In this context, services are aware of the interaction protocol that underlies their composition, and thus of the way in which they should interact. For this reason, choreography languages allow for and focus on the definition of protocols that parties should follow in order to achieve a common goal. There are two main approaches to define choreographies: (i) the *global model*, in which a protocol describes from a global perspective the messages exchanged by the parties, and (ii) the *interaction model* in which each service describes the temporal and logical dependencies among the messages it exchanges, i.e., a kind of interface definition. WS-CDL adopts the global model style, while WSCI and abstract processes of WS-BPEL are instances of the interaction model.

A common aspect considered by both orchestration and choreography styles is related to LRTs, i.e., the possibility of executing some parts of a composed service atomically. Nevertheless, atomicity does not imply here the usual "all-or-nothing" property of database transactions, because perfect rollback is unlikely in case of a fault. For example, the sending of a message cannot be undone. Consequently, LRTs often rely on a weaker notion of atomicity based on *compensations*. Compensations are ad hoc, user-programmed activities to be run when recovering from partial executions of LRTs arising after a fault or interruption because successful completion is no longer possible. For example, if some information has been sent and it is not longer valid after the fault, then a second message can be sent to the recipient.

Since, most industrial standards lack rigorous foundations, many efforts have been spent to provide a formal basis to reason about LRTs in composition languages. As far as the orchestration of LRT is concerned, the first proposal that appeared in the literature is (to the best of our knowledge) StAC (Butler and Ferreira 2004; Butler *et al.* 2002), which enriches an imperative language with primitives for installing, activating and removing compensations. After StAC, proposals such as (Bruni *et al.* 2005; Bruni *et al.* 2012; Butler *et al.* 2005b) have provided formal semantics for compensation languages, whose primitives are closer to real orchestration languages (see e.g., (Eisentraut and Spieler 2009)).

A different line of research focuses on the formal definition of LRTs for interaction-based choreographies (Bocchi *et al.* 2003; Bruni *et al.* 2004; Caires *et al.* 2009; de Vries *et al.* 2010; Laneve and Zavattaro 2005; Lucchi and Mazzara 2004). Typically, such research thread consists of extending well-assessed mobile calculi with ad hoc constructs tailored

to transactions and compensations. The enriched calculi exploit the communication primitives provided by, e.g., $\pi$-calculus (Milner *et al.* 1992) and Join calculus (Fournet *et al.* 1996), to model communication among parties. Hence, a composed service is described by a set of processes, any of them defining a particular partner of the complete system. In this way, any party declares the interface for proper composition with other partners. In this respect, transactional processes in those calculi are the formal counterpart of WSCI interfaces or BPEL abstract processes. Nevertheless, transactional calculi go beyond the scope of being just declarative definitions of service interfaces. In fact, they are aimed at providing an operational characterisation for business processes.

Consider the typical scenario in which a user books a room through a hotel reservation service. The ideal protocol followed by the two parties can be sketched as follow, by using an informal $\pi$-calculus-like notation:

$$\text{Client} \equiv request\,!(data).offer\,?(price).accept\,!(cc)$$
$$\text{Hotel} \equiv request\,?(details).offer\,!(rate).accept\,?(card).$$

We write $a!(v)$ for the sending of the message $v$ on the communication channel $a$, and $a?(x)$ for receiving on the variable $x$ some message sent on the channel $a$. The prefix symbol "$\_.\_$" must not be confused with the usual dot-notation from object-oriented language: it is used to establish the order in which actions must be executed. The client starts by sending a booking request to the hotel, which answers it with a rate offer. After receiving the offer, the client accepts it. This is the ideal protocol both parties should follow in order to accomplish the common goal. Nevertheless, there are several situations in which parties may be forced/inclined not to complete the execution of the protocol (e.g., the hotel has no available rooms for the requested day, or the client does not obtain acceptable rates). Clearly, just stopping the execution of the protocol may be not acceptable in most of the cases. Compensable transactions are designed to handle this kind of situations. In addition to the usual primitives of name passing calculi, transactional calculi provide a new kind of terms, generally of the form $[P : Q]$, involving: (i) a process $P$ that is required to be executed until completion and (ii) the corresponding compensation $Q$ to be executed in case $P$ cannot complete successfully. Moreover, the cancellation of the transaction can be handled by making $P$ reach a special process, usually denoted by *abort*. For example, when the hotel is unable to proceed with the order, it may abort the transaction and use the compensation to suggest an alternative hotel to the client (e.g., by sending the message $alt!(hotel)$). Then, the description of the protocol could be improved as follows:

$$\text{Client} \equiv [request\,!(data).offer\,?(price).(accept\,!(cc) + abort) : alt\,?(h).Q]$$
$$\text{Hotel} \equiv [request\,?(details).(offer\,!(rate).accept\,?(card) + abort) : alt\,!(hotel)].$$

The above protocol also allows the client to abort the transaction after receiving an offer (for instance, when the offer does not satisfy her expectations). Alternatively, the hotel may abort after receiving a request (for instance, when no rooms are available). Clearly, more sophisticated protocols may be written to allow clients and hotels to abort at any moment. We illustrate the use of compensations by making the component Hotel generate the single message $alt!(hotel)$ to provide the client with the information of an alternative hotel to contact running $Q$. More concretely, *hotel* could be a tuple of channel

names $request'$, $offer'$, $accept'$ to contact another hotel and $Q$ could be just a recursive instance of Client parametric on such channels.

Proposals in the literature differ mainly in the kinds of interactions allowed across transaction boundaries and the effects associated with the handshaking. Roughly, on one side of the design option we have completely permeable transactional scopes (as in $\pi$t-calculus (Bocchi *et al.* 2003)), where messages may freely cross transactional boundaries. For instance, a possible computation for the protocol can be described as follows:

$$
\begin{aligned}
\mathsf{Client}|\mathsf{Hotel} \;\to\; & [\mathit{offer}\,?(price).(\mathit{accept}\,!(cc) + abort) : \mathit{alt}\,?(h).Q] \\
& \qquad\qquad | \, [\mathit{offer}\,!(rate).\mathit{accept}\,?(card) + abort : \mathit{alt}\,!(hotel)] \\
\to\; & [(\mathit{accept}\,!(cc) + abort) : \mathit{alt}\,?(h).Q] \mid [\mathit{accept}\,?(card) : \mathit{alt}\,!(hotel)] \\
\to\; & [0 : \mathit{alt}\,?(h).Q] \mid [0 : \mathit{alt}\,!(hotel)] \\
\to\; & 0.
\end{aligned}
$$

Messages in the above computation flows freely from one transaction to the others. The main drawback of such approach is that transactional scopes does not ensure all interacting transactions to reach the same result, i.e., some of them may commit even though others have failed. For instance, consider a client executing the following protocol:

$$\mathsf{Client}' \equiv [\mathit{request}\,!(data).\mathit{offer}\,?(price).(\mathit{accept}\,!(cc)|abort) : \mathit{alt}\,?(h).Q].$$

Analogously to the previous case, there is a computation leading the system to the following state:

$$\mathsf{Client}'|\mathsf{Hotel} \to^* [(\mathit{accept}\,!(cc)|abort) : \mathit{alt}\,?(h).Q] \mid [\mathit{accept}\,?(card) : \mathit{alt}\,!(hotel)].$$

Then, the system may evolve to:

$$\mathsf{Client}'|\mathsf{Hotel} \to^* [abort : \mathit{alt}\,?(h).Q] \mid [0 : \mathit{alt}\,!(hotel)].$$

At this point, one transaction (that one from the client) can only abort by releasing the compensation $Q$, while the other (the hotel party) can only commit. Hence, the hotel has reserved a room that the client is not willing to book. It is true that we can write a different compensation for the client that contains the code needed for making the hotel cancel the reservation. For example, graceful termination mechanisms for closing dyadic sessions have been studied in (Boreale *et al.* 2008) and can be likely reused for transactions. Nevertheless, from our point of view, the fact that involved participants have no guaranties about the final outcome of the remaining transactional participants provides too weak a transactional mechanism for handling many common situations. Although, stronger transactional properties may be ensured by programming ad hoc coordination code through compensations, suitable transactional primitives should relieve programmers from writing such kind of code.

In this paper, we present *Committed Join* (cJoin), a calculus designed to ensure that all participants of the same transaction reach the same agreed outcome. This is achieved by making interacting transactional processes become part of the same larger transaction. The cJoin is an extension of the Join calculus (Fournet and Gonthier 1996), which is a process calculus with asynchronous name-passing communication. We based our approach on the Join calculus rather than on other more popular process calculi, such as the $\pi$-calculus,

because Join adheres to a locality principle that guarantees that extruded names cannot be used in input by the process who received them (they can only output values on such ports). This feature is crucial for deploying distributed implementations and it is not enforced in the full $\pi$-calculus. Moreover, it allows to obtain precise characterisations of transaction termination and atomic joining of multiple transactions, which are missing from most alternative proposals in the literature. Another advantage is that the operational semantics rules are quite simple and compact when compared to other transactional calculi.

The process Client′|Hotel behaves in cJoin as follows: when both transactions communicate through the port *request* for the first time they are merged in a unique larger transaction, whose transactional process and compensation correspond respectively to the parallel composition of the residuals of the original transactions and to the parallel composition of the original compensations, as shown below:

$$\text{Client}'|\text{Hotel} \rightarrow [\textit{offer}\,?(price).(\textit{accept}\,!(cc) \mid \textit{abort}\,) \\ \mid (\textit{offer}\,!(rate).\textit{accept}\,?(card) + \textit{abort}\,) : \textit{alt}\,?(h).Q \mid \textit{alt}\,!(hotel)].$$

From this moment on, the system may evolve as usual. In particular, assuming the hotel sends an offer (1) and the client sends the confirmation (2), the system moves as follows:

$$\rightarrow [\textit{accept}\,!(cc) \mid \textit{abort} \mid \textit{accept}\,?(card) : \textit{alt}\,?(h).Q \mid \textit{alt}\,!(hotel)] \quad (1)$$
$$\rightarrow [\textit{abort}\, : \textit{alt}\,?(h).Q \mid \textit{alt}\,!(hotel)]. \quad\quad\quad\quad\quad (2)$$

In this case, both original transactions are bound together and none of them has already committed, therefore, the abort condition reached by the client causes the hotel transaction to be compensated as well. In this way, transactional scopes of cJoin ensure that all parties of a transaction commit (resp. abort) only when all other parties commit (resp. abort), although each party is responsible for defining its own compensation. Note that the transactional primitive in cJoin relieves programmers from coding protocols needed to agree on a common result for a distributed transaction, while leaving to the programmer the responsibility for defining suitable compensations to recover aborted transactions. Though no automatic rollback mechanism is provided, it is obvious that restoring the initial process upon the abort can be straightforwardly programmed by recursive definitions like $P \equiv [Q : P]$, easy to implement in cJoin syntax.

Another important issue addressed in this paper is *transaction serializability*. Not to be confused with object serializability, it is a way for ensuring the correctness of reasoning at different levels of abstractions, in which transactions become atomic reductions when seen at the abstract level. Let us consider a set of $n$ transactions $\{T_i \mid 1 \leqslant i \leqslant n\}$, each consisting of several activities to be carried out. Their concurrent execution $T_1 \parallel \ldots \parallel T_n$ can interleave the activities from different transactions and it is said to be *serializable* if there exists a sequence $T_{i_1}, T_{i_2}, \ldots, T_{i_n}$ that executes all transactions one at a time (without interleaving their steps) and produces the same result (Bernstein *et al.* 1987). More generally, in the case of nested transactions, each $T_i$ could involve recursively several sub-transactions $T_{i,1}, \ldots, T_{i,n_1}$ among the activities to be carried out, whose execution is possibly interleaved with those of other transactions and of their sub-transactions. Serializability is important because it allows us to reason about the behaviour of a system by considering one transaction at a time, at any given level of nesting. Transaction

serializability is generally difficult to achieve in other proposals where communication is allowed across transactions. Here we show that, for a large class of cJoin processes, called *shallow processes*, serializability is guaranteed by construction because if two separately initiated transactions interact, then their scopes are merged together as part of the same transaction, i.e., after merge they cannot commit or abort independently.

For the prototype implementation of cJoin, we rely on available distributed implementation of Join. In fact, the primitives of Join have been exploited in the design of JoCaml (Conchon and Le Fessant 1999), an extension of the *Objective Caml*, a functional language with support of object-oriented and imperative paradigms, and Polyphonic C# (Benton *et al.* 2002) (later Cω) that extends C# with asynchronous methods and synchronisation patterns, called *chords*. We take advantage of this fact for extending JoCaml with transactional primitives. The resulting language, called *transactional* JoCaml (t-JoCaml), adds to JoCaml the possibility of writing programs that should execute as compensable transactions in the style of cJoin transactions.

*Paper Outline.* After introducing some preliminaries (Section 2), we give the syntax and semantics of cJoin (Section 3) and describe several examples illustrating the main features of cJoin (Section 4). By exploiting the strategy used for implementing cJoin transactions in Join itself, as summarised in Section 5; in Section 6, we describe t-JoCaml as an extension of JoCaml. We remark that t-JoCaml actually implements a sub-calculus of cJoin (called *flat*) in which transactions cannot be nested (see Section 3.2). Section 6.2.2 describes the corresponding extension of the JoCaml compiler we have realised. Section 7 shows how transaction serializability can be achieved in cJoin. To conclude, we compare our proposal against several approaches appeared in the literature and we present some final remarks (Section 8).

Preliminary studies on cJoin have been presented at IFIP-TCS 2004–IFIP 18th World Computer Congress, 3rd International Conference on Theoretical Computer Science, and COMETA 2003-workshop of the COMETA Project on Computational Metamodels, after which several other proposals of transactional process calculi emerged in the literature. Yet the features of cJoin remained quite peculiar and this work integrates previous studies with new perspectives in the area of service-oriented programming and business processes, most notably the well-disciplined use of compensations.

## 2. Preliminaries

In this section, we report on the operational semantics of the Join calculus as a chemical abstract machine, by following the presentation of (Fournet and Gonthier 1996).

### 2.1. *The chemical abstract machine*

The semantics of the Join calculus relies on the *reflexive chemical abstract machine* (CHAM). In a CHAM (Berry and Boudol 1992) computation states $S$ (called *solutions*) are finite multisets of terms $m$ (called *molecules*), and computations are multiset rewrites. Multisets are denoted by $m_1, \ldots, m_n$ and abbreviated with $\oplus_i m_i$. Solutions can be structured in a hierarchical way by using the operator *membrane* $\{[\cdot]\}$ to group a solution $S$ into a

(REACTION LAW)

$$\frac{m_1, \ldots, m_k \to m'_1, \ldots, m'_l \ \in \ \text{set of CHAM rules}}{m_1\sigma, \ldots, m_k\sigma \to m'_1\sigma, \ldots, m'_l\sigma}$$

(CHEMICAL LAW)

$$\frac{S \to S'}{S, S'' \to S', S''}$$

(MEMBRANE LAW)

$$\frac{S \to S'}{\{[S]\} \to \{[S']\}}$$

Fig. 1. CHAM laws.

molecule $\{[S]\}$. (In (Berry and Boudol 1992) molecules can be built also with the constructor *airlock*, but it is not needed in our presentation.)

Transformations are described by a set of *chemical rules*, which specify how solutions react. In a CHAM, there are two different kinds of chemical rules: *heating/cooling* rules $S \leftrightharpoons S'$ representing syntactical rearrangements of molecules in a solution, and *reaction* rules $S \to S'$. Heating/cooling rules are analogous to the axioms for structural congruence in process calculi, and thus called also *structural* rules. Structural rules are reversible: a solution obtained by applying a cooling rule can be heated back to the original state, and vice versa. Reaction rules, on the other hand, cannot be undone. Rules can carry formal parameters to be matched against actual parameters in the redex and substituted in the right-hand side.

The laws governing CHAM computations are in Figure 1 (we give them for reaction rules, but they are applicable to heating/cooling rules as well):

— *Reaction law*: given a rule, an instance of its left-hand side can be replaced by the corresponding instance of the right-hand side. The substitution $\sigma$ replaces the formal parameters with the actual parameters by matching the solution against the left-hand side of the rule.
— *Chemical law*: reactions can be applied in every larger solution.
— *Membrane law*: reactions may occur at any level in the hierarchy of solutions.

Note that CHAM's heating/cooling/reaction rules have no premises and are purely local. They specify only the part of the solution that actually changes. Moreover, since solutions are multisets, not overlapping rules can be applied concurrently.

### 2.2. *The Join calculus*

The Join calculus relies on an infinite set of names $x, y, u, v, \ldots$. Name tuples are written $\check{u}$. Join processes, definitions and patterns are in Figure 2(a). A *process* is either the inert process 0, the asynchronous emission $x\langle\check{y}\rangle$ of message $\check{y}$ on port $x$, the process **def** $D$ **in** $P$ equipped with local ports defined by $D$, or a parallel composition of processes $P|Q$. A *definition* is a conjunction of elementary reactions $J \triangleright P$ that associate *join-patterns* $J$ with *guarded processes* $P$. Names defined by $D$ in **def** $D$ **in** $P$ are bound in $P$ and in all the guarded processes contained in $D$. The sets of defined names *dn*, received names *rn* and free names *fn* are in Figure 2(b).

**Example 1.** Consider the processes $Q = $ **def** $proxy\langle y\rangle \triangleright server\langle proxy, y\rangle$ **in** $proxy\langle a\rangle$ and $P = server\langle proxy, b\rangle|Q$. Roughly, $Q$ defines a local port *proxy* such that when a message on *proxy* arrives with any content $y$ then the name *proxy* is extruded on (the elsewhere defined, free port) *server* together with $y$. Intuitively, the "local proxy" forwards to the "public

$$
\begin{array}{lll}
(\text{PROC}) & P, Q & ::= \quad 0 \mid x\langle \vec{y} \rangle \mid \textbf{def } D \textbf{ in } P \mid P|Q \\
(\text{DEF}) & D, E & ::= \quad J \triangleright P \mid D \wedge E \\
(\text{PAT}) & J, K & ::= \quad x\langle \vec{y} \rangle \mid J|K
\end{array}
$$

(a) Syntax

$$
(\text{FREE})
$$

$$
\begin{array}{llllll}
fn(0) & = & \emptyset & fn(x\langle \vec{y} \rangle) & = & \{x\} \cup \{\vec{y}\} \\
fn(\textbf{def } D \textbf{ in } P) & = & (fn(P) \cup fn(D)) \backslash dn(D) & fn(P|Q) & = & fn(P) \cup fn(Q) \\
fn(J \triangleright P) & = & dn(J) \cup (fn(P) \backslash rn(J)) & fn(D \wedge E) & = & fn(D) \cup fn(E)
\end{array}
$$

$$
(\text{DEFINED})
$$

$$
\begin{array}{llllll}
dn(J \triangleright P) & = & dn(J) & dn(D \wedge E) & = & dn(D) \cup dn(E) \\
dn(x\langle \vec{y} \rangle) & = & \{x\} & dn(J|K) & = & dn(J) \cup dn(K)
\end{array}
$$

$$
(\text{RECEIVED})
$$

$$
\begin{array}{llllll}
rn(x\langle \vec{y} \rangle) & = & \{\vec{y}\} & rn(J|K) & = & rn(J) \cup rn(K)
\end{array}
$$

(b) Free, Defined and Received names

$$
\begin{array}{lll}
(\text{STR-NULL}) & 0 \rightleftharpoons \\
(\text{STR-JOIN}) & P \mid Q \rightleftharpoons P, Q \\
(\text{STR-AND}) & D \wedge E \rightleftharpoons D, E \\
(\text{STR-DEF}) & \textbf{def } D \textbf{ in } P \rightleftharpoons D\sigma_{dn(D)}, P\sigma_{dn(D)} & (range(\sigma_{dn(D)}) \text{ globally fresh}) \\
(\text{RED}) & J \triangleright P, J\sigma \rightarrow J \triangleright P, P\sigma
\end{array}
$$

(c) Semantics

Fig. 2. Join calculus.

server" each message tagged with its "origin" (i.e., the name of the proxy). The initial state of $Q$ carries a message on *proxy* whose content is $a$. The process $P$ places $Q$ in a context that includes a message to port *server* with content $\langle proxy, b \rangle$. Then, $fn(Q) = \{server, a\}$ and $fn(P) = \{server, a, b, proxy\}$. Therefore, the name *server* has common meaning in $P$ and $Q$, while the symbol *proxy* denotes, by accident, different ports in $server\langle proxy, b \rangle$ and $Q$: a free (elsewhere defined) port in the former and a private port in the latter. Moreover, letting $D = proxy\langle y \rangle \triangleright server\langle proxy, y \rangle$, we have $fn(D) = \{proxy, server\}$, $dn(D) = \{proxy\}$ and $rn(D) = \{y\}$.

The semantics of the Join calculus relies on the *reflexive* CHAM. It is called reflexive because active reaction rules are represented by molecules present in solutions, which are activated dynamically. Molecules, generated by $m ::= P \mid D$, correspond to terms of the Join calculus denoting processes or definitions. The chemical rules are shown in Figure 2(c). Rule STR-NULL states that 0 can be added or removed from any solution. Rules STR-JOIN and STR-AND implies the associativity and commutativity of $|$ and $\wedge$, because $\_,\_$ is such. STR-DEF denotes the activation of a local definition, which implements a static scoping discipline by properly renaming defined ports by *globally fresh* names. A name $x$ is *fresh* w.r.t. a process $P$ (resp. a definition $D$) if $x \notin fn(P)$ (resp. $x \notin fn(D)$). Moreover, $x$ is fresh w.r.t. a solution $s$ if it is fresh w.r.t. every term in $s$. A set of names $X$ is fresh if every name in $X$ is such. We write the substitution of names $x_1, \ldots, x_n$ by names $y_1, \ldots, y_n$

as $\sigma = \{^{y_1,\ldots,y_n}/_{x_1,\ldots,x_n}\}$, with $\mathrm{dom}(\sigma) = \{x_1,\ldots,x_n\}$ and $\mathrm{range}(\sigma) = \{y_1,\ldots,y_n\}$. We indicate with $\sigma_N$ an injective substitution $\sigma$ such that $\mathrm{dom}(\sigma) = N$. When we require names to be globally fresh, we mean that they must be different from all other names appearing in the enclosing context.

**Example 2.** Consider the process $P$ in the previous example. A solution $s$ containing only $P$, i.e., $s = \{[P]\}$, may be heated as follows. First, the two parallel agents are separated by obtaining $\{[server\langle proxy, b\rangle, \mathbf{def}\ proxy\langle y\rangle \rhd server\langle proxy, y\rangle\ \mathbf{in}\ proxy\langle a\rangle]\}$. Now, the second molecule contains a definition of a local port *proxy* that is different from the homonymous free port *proxy* in the first molecule $server\langle proxy, b\rangle$. Hence, when using STR-DEF for separating the local definition from the corresponding process, the local definition of *proxy* must be renamed by using a fresh name, say $proxy_1$, obtaining the solution $s' = \{[server\langle proxy, b\rangle, proxy_1\langle y\rangle \rhd server\langle proxy_1, y\rangle, proxy_1\langle a\rangle]\}$.

Finally, RED describes the use of an active reaction rule ($J \rhd P$) to consume messages forming an instance of $J$ (for a suitable substitution $\sigma$, with $dom(\sigma) = rn(J)$), and produce a new instance $P\sigma$ of its guarded process $P$.

**Example 3.** By applying rule (RED) to the solution $s'$ from Example 2 for $\sigma = \{^a/_y\}$, we get $s' \to \{[server\langle proxy, b\rangle, proxy_1\langle y\rangle \rhd server\langle proxy_1, y\rangle, server\langle proxy_1, a\rangle]\}$. Note that the local port $proxy_1$ has been extruded on the free channel *server*.

**Remark 2.1.** For $\pi$-calculus enthusiasts, the Join calculus can be easily grasped by considering the main differences enforced by the syntax, namely: (i) like in the asynchronous $\pi$-calculus, only output particles are allowed, not output prefixes; (ii) input prefixes are encoded as definitions and joint inputs are allowed (i.e., more than one message can be consumed atomically); (iii) all definitions are persistent, as if they were prefixed by the replication operator; (iv) the definition construct $\mathbf{def}\ D\ \mathbf{in}\ P$ is also binding all defined names in $D$ to ensure the unique receptor property and favour distributed implementation (names can still be extruded, but the processes that receive them can only output on them); (v) the programming style is continuation passing, in the sense that output prefixes can be encoded by including in the message a fresh continuation name $k$ where the acknowledge of the receipt must be sent to activate the output-prefixed process. For example, the process $P$ from Example 1 can be understood as the $\pi$-calculus process $\bar{s}\langle p, b\rangle \mid (\nu p)(\ \bar{p}a \mid !p(y).\bar{s}\langle p, y\rangle\ )$, where for brevity initial letters of port names are used.

We shall write $P \equiv Q$ when $P \leftrightharpoons^* R$, for $\leftrightharpoons^*$ the reflexive and transitive closure of the relation $\leftrightharpoons$. Moreover, we abuse the notation by allowing one step reductions up to $\equiv$, i.e., writing $P \to Q$ when $P \leftrightharpoons^* \to \leftrightharpoons^* Q$. We write $P \to^n Q$ for $n \geqslant 0$ if there exist $n+1$ processes $P_0, \ldots, P_n$ such that $P \equiv P_0 \to P_1 \to \ldots \to P_{n-1} \to P_n \equiv Q$. Finally, we write $P \to^* Q$ if $P \to^n Q$ for some $n$.

Note that by exploiting heating and cooling rules, it is always possible to move definitions around, possibly after some renaming of their defined ports. In particular, we remark that $\mathbf{def}\ D\ \mathbf{in}\ P \equiv P|\mathbf{def}\ D\ \mathbf{in}\ 0$ whenever $fn(P) \cap dn(D) = \varnothing$ and furthermore that $(\mathbf{def}\ D_1\ \mathbf{in}\ P_1)|(\mathbf{def}\ D_2\ \mathbf{in}\ P_2) \equiv \mathbf{def}\ D_1\ \mathbf{in}\ (P_1|\mathbf{def}\ D_2\ \mathbf{in}\ P_2) \equiv \mathbf{def}\ D_1 \wedge D_2\ \mathbf{in}\ (P_1|P_2)$ if $dn(D_1) \cap (dn(D_2) \cup fn(P_2)) = dn(D_2) \cap (dn(D_1) \cup fn(P_1)) = \varnothing$.

$$
\begin{array}{llll}
(\textsc{mess}) & M, N & ::= & 0 \mid x\langle\vec{y}\rangle \mid M|N \\
(\textsc{proc}) & P, Q & ::= & M \mid \mathbf{def}\ D\ \mathbf{in}\ P \mid P|Q \mid [P:Q] \mid abort \\
(\textsc{def}) & D, E & ::= & J \triangleright P \mid D \wedge E \mid J \blacktriangleright P \\
(\textsc{pat}) & J, K & ::= & x\langle\vec{y}\rangle \mid J|K
\end{array}
$$

Fig. 3. Syntax of cJoin.

In Sections 3 and 4, where several examples of cJoin processes are given, we shall use the following syntactic sugar for the case where different decisions can be taken after receiving a certain message tuple: we write $J \triangleright P + Q$ in place of $J \triangleright P \wedge J \triangleright Q$.

## 3. Committed join

In order to handle LRTs we extend the syntax of the Join calculus as shown in Figure 3. For convenience, we introduce the syntactical category $M$ of processes without definitions, i.e., a parallel composition of messages. In addition to cJoin processes, we add terms $[P : Q]$ to denote transactions, where $P$ is the transactional process and $Q$ is its compensation, i.e., the process to be executed when $P$ aborts. A transactional process is executed in isolation until reaching either a commit state or an abort condition. If $P$ commits, the obtained result is delivered to the outside of the transaction. Otherwise, the compensation $Q$ is activated. The abort decision is caused by the presence of the special basic process *abort*.

A new kind of definitions $J \blacktriangleright P$, called *merge definitions*, is introduced to specify the possible interactions among transactions (inter-transaction communication). Merge definitions allow the consumption of messages produced in the scope of different transactions by joining all participants in a unique larger LRT. Note that all ongoing transactions that want to merge are treated uniformly: they will issue a request to $J \blacktriangleright P$ by producing an instance of a particle in $J$ and when a full instance $J\sigma$ can be formed out of those particles all the transactions containing those particles are merged. Thus, $J \blacktriangleright P$ acts as some kind of *message board* where merge advertises are posted by transactions. This is different w.r.t. an "asymmetric" merge discipline, where one message in a transaction is received by an input running in a different transaction causing the merge. In fact, this kind of permeability of transaction scopes is not allowed in cJoin. Moreover, cJoin definitions can be used to create transactions dynamically. For instance, by firing $J \triangleright [P : Q]$ a new instance of the transaction $P$ with compensation $Q$ is activated. Perfect rollback can thus be programmed just by writing definitions like $J \triangleright [P : J]$. In fact, if an instance $J\sigma$ of $J$ is consumed to produce $[P\sigma : J\sigma]$ such that $P\sigma$ will abort, then the original messages $J\sigma$ are restored as a compensation. Notably, when two or more of such transactions are merged into a unique transaction, then the overall compensation becomes just the parallel composition of their perfect rollbacks.

The sets of defined names *dn*, received names *rn* and free names *fn* are defined in Figure 4. In particular, we distinguish between defined *ordinary* names $dn_{\mathrm{o}}(D)$ and defined *merge* names $dn_{\mathrm{m}}(D)$ that, as a general well-formedness discipline, are always assumed to be disjoint sets of names.

(PROC)
$$fn(0) = \emptyset$$
$$fn(P|Q) = fn(P) \cup fn(Q)$$
$$fn(\mathbf{def}\ D\ \mathbf{in}\ P) = (fn(P) \cup fn(D)) \backslash dn(D)$$

$$fn(x\langle\vec{y}\rangle) = \{x\} \cup \{\vec{y}\}$$
$$fn(abort) = \emptyset$$
$$fn([P:Q]) = fn(P) \cup fn(Q)$$

(DEF)
$$fn(J \rhd P) = dn(J) \cup (fn(P) \backslash rn(J))$$
$$fn(J \blacktriangleright P) = dn(J) \cup (fn(P) \backslash rn(J))$$

$$fn(D \wedge E) = fn(D) \cup fn(E)$$

$$dn_{\mathrm{o}}(J \rhd P) = dn(J)$$
$$dn_{\mathrm{o}}(J \blacktriangleright P) = \emptyset$$

$$dn_{\mathrm{o}}(D \wedge E) = dn_{\mathrm{o}}(D) \cup dn_{\mathrm{o}}(E)$$

$$dn_{\mathrm{m}}(J \rhd P) = \emptyset$$
$$dn_{\mathrm{m}}(J \blacktriangleright P) = dn(J)$$

$$dn_{\mathrm{m}}(D \wedge E) = dn_{\mathrm{m}}(D) \cup dn_{\mathrm{m}}(E)$$

(PAT)
$$rn(x\langle\vec{y}\rangle) = \{\vec{y}\}$$
$$dn(x\langle\vec{y}\rangle) = \{x\}$$

$$rn(J|K) = rn(J) \cup rn(K)$$
$$dn(J|K) = dn(J) \cup dn(K)$$

Fig. 4. Free, defined and received names.

| (STR-NULL) | $0$ | $\rightleftharpoons$ | |
|---|---|---|---|
| (STR-JOIN) | $P \mid Q$ | $\rightleftharpoons$ | $P, Q$ |
| (STR-AND) | $D \wedge E$ | $\rightleftharpoons$ | $D, E$ |
| (STR-DEF) | $\mathbf{def}\ D\ \mathbf{in}\ P$ | $\rightleftharpoons$ | $D\sigma_{dn(D)}, P\sigma_{dn(D)}$ |
| | | | $(range(\sigma_{dn(D)})\ globally\ fresh)$ |
| (RED) | $J \rhd P, J\sigma$ | $\rightarrow$ | $J \rhd P, P\sigma$ |
| (STR-TRANS) | $[P : Q]$ | $\rightleftharpoons$ | $\{\!\![P, \llcorner Q \lrcorner]\!\!\}$ |
| (COMMIT) | $\{\!\![M|\mathbf{def}\ D\ \mathbf{in}\ 0, \llcorner Q \lrcorner]\!\!\}$ | $\rightarrow$ | $M$ |
| (ABORT) | $\{\!\![abort|P, \llcorner Q \lrcorner]\!\!\}$ | $\rightarrow$ | $Q$ |
| (MERGE) | | | |
| $(\Pi_i J_i \blacktriangleright P), \otimes_i \{\!\![J_i\sigma, S_i, \llcorner Q_i \lrcorner]\!\!\}$ | | $\rightarrow$ | $(\Pi_i J_i \blacktriangleright P), \{\!\![\otimes_i S_i, P\sigma, \llcorner \Pi_i Q_i \lrcorner]\!\!\}$ |

Fig. 5. Operational semantics of cJoin.

## 3.1. *Operational semantics*

The operational semantics of cJoin is given in the reflexive CHAM style. Molecules $m$ and solutions $S$ for cJoin are as follows:

$$m ::= P \mid D \mid \llcorner P \lrcorner \mid \{\!\![S]\!\!\}$$
$$S ::= m \mid m, S.$$

Processes and definitions are molecules. Terms $\llcorner Q \lrcorner$ denote compensations, i.e., frozen processes to be activated only when the corresponding transaction aborts. Molecules $\{\!\![S]\!\!\}$ stands for running transactions.

The chemical rules for cJoin are given in Figure 5. The first five chemical rules are the ordinary ones for Join. Rule (STR-TRANS) states that a term denoting a transaction corresponds to a sub-solution consisting of two molecules: the transactional process $P$

and its compensation $Q$, which is frozen (the operator $\llcorner \lrcorner$ forbids the enclosed process to compute because there is no rule for applying reductions inside it).

A transaction can commit only when all internal computations have finished. This situation can be characterised as the special pattern $\{[M|\mathbf{def}\ D\ \mathbf{in}\ 0, \llcorner Q \lrcorner]\}$, comprising some messages $M$, the definition $D$ for computing inside the transaction and the so-far installed compensation $Q$. Since the scope of $D$ is the nil process, it is evident that at commit time, private definitions of the transactional process can be discarded, because neither the messages that are being released contain those names nor they could have been extruded previously. On the other hand, if defined names of $D$ were present in some of the messages inside the transaction, then the transactional activity would not yet be complete and the commit would not take place. When a transaction commit (rule (COMMIT)), the local resources $M$ produced inside a transaction are released as outcome. After commit, its compensation procedure $\llcorner Q \lrcorner$ is useless and it is discarded as well.

The abortion of a transaction is handled by the rule (ABORT), which releases $Q$ whenever *abort* is present in the solution.

Interactions among transactions are dealt with (MERGE), where we let the notation $\Pi_i J_i$ and $\Pi_i Q_i$ denote, respectively, the parallel composition of messages $J_1 | \ldots | J_n$ and the parallel composition of processes $Q_1 | \ldots | Q_n$. We also recall that the notation $\otimes_i m_i$ denotes multisets of molecules and solutions. The rule (MERGE) consumes messages from different transactions and creates a larger transaction by combining the definitions and messages of the original ones with a new instance of the guarded process $P\sigma$, where $dom(\sigma) = rn(J_1 | \ldots | J_n)$. Name clashes are avoided because we assume that STR-DEF generates globally fresh names. The compensation for the joint transaction is the parallel composition of all the original compensations.

The rule (MERGE) is quite general, as it can be used to join atomically an unbounded number of ongoing transactions. To help the understanding, we show a few particular instances of the rule (MERGE), which will be used frequently in the rest of this paper.

The first case is that of a trivial merge, where a unique transaction is involved. It is interesting to show it just to clarify that rule (MERGE) can be applied without merging transactions. This kind of merge rules works as global definitions that can be used in any transaction of the system.

(TRIV-MERGE)
$$(x\langle y \rangle \blacktriangleright P), \{[x\langle v \rangle, S, \llcorner Q \lrcorner]\} \;\to\; (x\langle y \rangle \blacktriangleright P), \{[S, P\{^v/_y\}, \llcorner Q \lrcorner]\}.$$

The second case is that of a two-way merge, which will be very useful in our examples.

(TWO-MERGE)
$$
\begin{array}{l}
(x_1\langle y_1 \rangle | x_2\langle y_2 \rangle \blacktriangleright P), \\
\{[x_1\langle v_1 \rangle, S_1, \llcorner Q_1 \lrcorner]\}, \\
\{[x_2\langle v_2 \rangle, S_2, \llcorner Q_2 \lrcorner]\}
\end{array}
\;\to\; (x_1\langle y_1 \rangle | x_2\langle y_2 \rangle \blacktriangleright P), \{[S_1, S_2, P\{^{v_1,v_2}/_{y_1,y_2}\}, \llcorner Q_1 | Q_2 \lrcorner]\}.
$$

Note that, as a degenerate case of rule (TWO-MERGE) we also have:

(TWO-MERGE-DEG)
$$
\begin{array}{l}
(x_1\langle y_1 \rangle | x_2\langle y_2 \rangle \blacktriangleright P), \\
\{[x_1\langle v_1 \rangle, x_2\langle v_2 \rangle, S, \llcorner Q \lrcorner]\}
\end{array}
\;\to\; (x_1\langle y_1 \rangle | x_2\langle y_2 \rangle \blacktriangleright P), \{[S, P\{^{v_1,v_2}/_{y_1,y_2}\}, \llcorner Q \lrcorner]\}.
$$

$$
\begin{aligned}
\mathsf{HB} \quad &\equiv \quad \textbf{def} \;\; HotelSrv\langle r\rangle \mid HotelReq\langle d,\kappa\rangle \;\;\blacktriangleright\;\; r\langle d,\kappa\rangle \\
&\quad\;\; \textbf{in} \;\; \mathsf{H} \mid \mathsf{C} \\[6pt]
\mathsf{H} \quad &\equiv \quad [\textbf{def} \; request\langle details,\kappa\rangle \;\triangleright\; \kappa\langle price, accept\rangle + abort \\
&\qquad\quad \wedge \; accept\langle cc\rangle \;\triangleright\; 0 \\
&\qquad \textbf{in} \;\; HotelSrv\langle request\rangle : alt\langle hotel\rangle] \\[6pt]
\mathsf{C} \quad &\equiv \quad [\textbf{def} \; offer\langle rate,\kappa\rangle \;\triangleright\; \kappa\langle card\rangle + abort \\
&\qquad \textbf{in} \;\; HotelReq\langle data, offer\rangle : Q]
\end{aligned}
$$

Fig. 6. Hotel booking.

The following proposition states that Join is a sub-calculus of Join.

**Proposition 3.1.** Join is a sub-calculus of cJoin.

*Proof.* It is obvious from the syntax that any Join process is also a cJoin process. It remains to show that: (1) for any Join processes $P$ and $Q$, if $P \rightarrow Q$ in Join, then $P \rightarrow Q$ in cJoin; and (2) for any Join processes $P$, if $P \rightarrow Q$ in cJoin, then $Q$ is a Join process and $P \rightarrow Q$ in Join. Both implications follow straightforwardly from the fact that the chemical rules of cJoin can be partitioned in two sets: one consisting exactly of the chemical rules of Join, and the other containing structural rules involving non-Join operators on both sides and reaction rules involving non-Join operators in the left-hand side. □

We are now ready to revisit the hotel booking problem described in Section 1 and show how it can be modelled in cJoin.

**Example 4.** Process HB in Figure 6, shows a possible modelling for the hotel booking application, where H describes the behaviour of the hotel service while C models the protocol followed by the client. There are two main differences with the description given in Section 1. First, we adhere to the continuation-passing style for enabling communication among different processes. Note that channels *request* and *offer* carry on one extra parameter (the continuation) that identifies the channel $\kappa$ where to communicate next. Second, the system is not just the parallel composition of the two parties, but it also contains a merge definition that allows the communication among the two transactions. In fact, the two parties do not start by communicating directly and the first interaction takes place indirectly. Note that C starts by sending the message $HotelReq\langle data, offer\rangle$ to a merge channel. Similarly, H sends $HotelSrv\langle request\rangle$. While neither C nor H can complete their transactions in isolation, these two messages together enable the merge definition that forwards the request from the client to the hotel and joins both transactions, producing

the following state:

$$HB \equiv \textbf{def} \quad HotelSrv\langle r\rangle \mid HotelReq\langle d,\kappa\rangle \blacktriangleright r\langle d,\kappa\rangle$$

$$\textbf{in} \; [\textbf{def} \;\; request\langle details,\kappa\rangle \rhd \kappa\langle price, accept\rangle \tag{1}$$
$$+\; abort \tag{2}$$
$$\wedge \;\; accept\langle cc\rangle \rhd 0 \tag{3}$$
$$\wedge \;\; offer\langle rate,\kappa\rangle \rhd \kappa\langle card\rangle \tag{4}$$
$$+\; abort \tag{5}$$
$$\textbf{in} \;\; request\langle data, offer\rangle : alt\langle hotel\rangle \mid Q].$$

At this moment the hotel receives the request from the client and it can (nondeterministically) decide whether to make an offer (reaction (1)) or to abort (reaction (2)). We remind that the notation $J \rhd P + Q$ is syntactic sugar for $J \rhd P \wedge J \rhd Q$. If the hotel aborts, then both transactions are aborted and the corresponding compensations (i.e., $alt\langle hotel\rangle \mid Q$) are activated. Otherwise, the hotel can produce the message $offer\langle price, accept\rangle$ to send an offer to the client who, in turn, may decide whether to accept it (reaction (4)) or to abort the conversation (rule (5)). The abortion from the client is handled analogously to the previous case. If the client accepts the offer, then it generates the message $accept\langle card\rangle$, which enables the reaction rule (3). When reaction (3) is fired all messages inside the transaction are consumed, and hence the transaction commits (i.e., both parties have successfully finished).

Another interesting example is the modelling of a mailing-list manager with all-or-nothing delivery of messages to subscribers.

**Example 5 (Mailing list).** Consider a data structure that allows to send atomically a message to a list of subscribers (in the sense that the same message is either sent to all or to none). Such structure can be defined as $ML = MailingList\langle k\rangle \rhd MLDef$, where:

$$MLDef \equiv \textbf{def} \; List \; \textbf{in} \; k\langle add, tell, close\rangle \mid l\langle nil\rangle$$

$$List \equiv \;\; nil\langle v,w\rangle \blacktriangleright w\langle\rangle$$
$$\wedge \;\; l\langle y\rangle \mid add\langle x\rangle \rhd \textbf{def} \; z\langle v,w\rangle \blacktriangleright x\langle v\rangle \mid y\langle v,w\rangle \; \textbf{in} \; l\langle z\rangle$$
$$\wedge \;\; l\langle y\rangle \mid tell\langle v\rangle \rhd [\textbf{def} \; z\langle\rangle \rhd 0 \; \textbf{in} \; y\langle v,z\rangle \mid l\langle y\rangle : l\langle y\rangle]$$
$$\wedge \;\; l\langle y\rangle \mid close\langle\rangle \rhd 0.$$

A new mailing list is created by sending a message to the port *MailingList*. Since, cJoin adheres to the "continuation passing" style of programming, the content of the message sent to *MailingList* is a continuation port $k$, which expects information about the newly created mailing list. The creation of a new list defines five fresh ports *nil, l, add, tell* and *close*: three of them (namely *add*, *tell*, and *close*) will be used to interact with the list from "outside" and will be sent to the port $k$ as the outcome of the creation. The remaining two ports will never be extruded. They denote the empty list (*nil*) and the actual state of the list (*l*). Once a list is created, a new subscriber can be added by sending a message *add* with the name $x$ of the port where it will be listening to for new messages. In this case, the list is modified by installing $z$ (on top of it), a forwarder of messages to $x$.

The port *tell* is used to send a message $v$ to the list. When *tell* is received a new transaction identified by a fresh name $z$ is generated, and the state of the structure is put

inside the transaction, therefore all other activities, such as adding or closing are blocked until the transaction ends. Inside the transaction, the message $v$ is sent to the forwarder at the top of the list $y$ with the identifier of the transaction $z$. Note that each forwarder sends the message to the corresponding subscriber and to the following forwarder in the list. This is repeated until *nil* is reached, when a message to the identifier of the transaction is sent. The firing rule $z\langle\rangle \triangleright 0$ consumes the last local name and the transaction commits by releasing all the messages addressed to the subscribers and the state of the list. Then the list is ready to serve new requests. The following process Sys subscribes two users, *Alice* and *Bob* to the mailing list, and sends the message *News*.

$$\mathsf{Emp} \equiv employees\langle a, t, c\rangle \triangleright a\langle Alice\rangle \mid a\langle Bob\rangle \mid t\langle News\rangle$$

$$\mathsf{Sys} \equiv \mathbf{def}\ \mathsf{ML} \wedge \mathsf{Emp}\ \mathbf{in}\ MailingList\langle employees\rangle$$

A possible computation of the process Sys is shown in Figure 7. In this particular computation, both subscriptions take place before the emission of the message, nevertheless the process does not fix this priority and consequently messages could be consumed in a different order. For simplicity, we abbreviate chemical solutions by omitting definitions present in successive solutions, though we usually write only those definitions involved in the reduction step. Inside solutions, we underline the fired reaction rule and the consumed messages that matched the corresponding pattern. The phase LIST CREATION instantiates a new mailing list by defining the fresh ports *add*, *tell* and *close*, which are sent to the port *employee*. The second phase (SUBSCRIPTIONS) adds the names *Alice* and *Bob* to the created mailing list. Phase DISTRIBUTION OF *News* generates a new transaction that produces (in a sequential way) the copies of the name *News* to be sent to any subscriber of the list. Nevertheless, those messages are not released until the phase COMMIT takes place. Only when the transaction commits, the generated messages are atomically sent to subscribers.

## 3.2. *Flat transactions*

Nesting is a useful mechanism for programming transactions, e.g., when a transaction can succeed even when certain sub-activities fail. In the area of databases, nested transactions have been studied since (Moss 1981). Contrary to other languages proposed in the literature that do not support nesting (e.g., Web$\pi_\infty$ (Lucchi and Mazzara 2004) and $\rho\pi$ (Lanese *et al.* 2010a)), cJoin syntax allows for proper nested transactions, like in $[[P : P'] : [Q : Q']]$. Nevertheless, many common situations that would involve nested transactions can be modelled in cJoin without nesting by exploiting dynamic merge and message-passing communication, as shown in the multi-way transaction example (see Section 4.2).

This section introduces *flat* cJoin, which is a sub-calculus of cJoin without nested transactions. The following sections will show that flat cJoin is expressive enough for modelling several common programming patterns (Section 4) and, in addition, how it can be implemented (Sections 5 and 6). In fact, as it will be clear later, the syntactic restrictions imposed on flat cJoin help us to encode flat cJoin back to Join and to extend existing distributed implementation of Join to implement flat cJoin. Flat cJoin is defined

Initial Soup: $\{\!\{\,\underline{\text{Sys}}\,\}\!\} \rightleftharpoons \{\!\{\,[MailingList\langle k\rangle \triangleright \text{MLDef}, \text{Emp}, MailingList\langle employees\rangle\,\}\!\}$

SMALL CAPS: LIST CREATION:

$\{\!\{\,[\underline{MailingList\langle k\rangle \triangleright \text{MLDef}}, \text{Emp}, \underline{MailingList\langle employees\rangle}\,\}\!\} \rightarrow$

$\{\!\{[MailingList\langle k\rangle \triangleright \text{MLDef}, \text{Emp}, \underline{\text{MLDef}\{{}^{employees}/_k\}}\,\}\!\} \rightleftharpoons$

$\{\!\{[MailingList\langle k\rangle \triangleright \text{MLDef}, \underline{\text{Emp}}, \text{List}\{{}^{employees}/_k\}, \underline{employees\langle add, tell, close\rangle}, l\langle nil\rangle]\,\}\!\} \rightarrow$

$\{\!\{l\langle y\rangle \mid add\langle x\rangle \triangleright \dots, \dots, l\langle nil\rangle, add\langle Alice\rangle, add\langle Bob\rangle, tell\langle News\rangle]\,\}\!\}$

SUBSCRIPTIONS:

$\{\!\{l\langle y\rangle \mid add\langle x\rangle \triangleright \dots, \dots, \underline{l\langle nil\rangle}, \underline{add\langle Alice\rangle}, add\langle Bob\rangle, tell\langle News\rangle]\,\}\!\} \rightarrow$

$\{\!\{l\langle y\rangle \mid add\langle x\rangle \triangleright \dots, \dots, z_A\langle v, w\rangle \blacktriangleright Alice\langle v\rangle|nil\langle v, w\rangle, \underline{l\langle z_A\rangle}, \underline{add\langle Bob\rangle}, tell\langle News\rangle]\,\}\!\} \rightarrow$

$\{\!\{l\langle y\rangle \mid tell\langle v\rangle \triangleright \dots, \dots, z_B\langle v, w\rangle \blacktriangleright Bob\langle v\rangle|z_A\langle v, w\rangle, l\langle z_B\rangle, tell\langle News\rangle]\,\}\!\}$

DISTRIBUTION OF *News*:

$\{\!\{l\langle y\rangle \mid tell\langle v\rangle \triangleright \dots, \dots, z_B\langle v, w\rangle \blacktriangleright Bob\langle v\rangle|z_A\langle v, w\rangle, \underline{l\langle z_B\rangle}, \underline{tell\langle News\rangle}]\,\}\!\} \rightarrow$

$\{\!\{z_B\langle v, w\rangle \blacktriangleright Bob\langle v\rangle|z_A\langle v, w\rangle, \dots, [\underline{z\langle\rangle \triangleright 0}, \underline{z_B\langle News, z\rangle}, l\langle z_B\rangle, \llcorner l\langle z_B\rangle \lrcorner]\,\}\!\} \rightarrow$

$\{\!\{z_A\langle v, w\rangle \blacktriangleright Alice\langle v\rangle|nil\langle v, w\rangle, \dots, [z\langle\rangle \triangleright 0, Bob\langle News\rangle, \underline{z_A\langle News, z\rangle}, l\langle z_B\rangle, \llcorner l\langle z_B\rangle \lrcorner]\,\}\!\} \rightarrow$

$\{\!\{[nil\langle v, w\rangle \blacktriangleright w\langle\rangle, \dots, [z\langle\rangle \triangleright 0, Bob\langle News\rangle, Alice\langle News\rangle, \underline{nil\langle News, z\rangle}, l\langle z_B\rangle, \llcorner l\langle z_B\rangle \lrcorner]\,\}\!\} \rightarrow$

$\{\!\{\dots, [\underline{z\langle\rangle \triangleright 0}, Bob\langle News\rangle, Alice\langle News\rangle, \underline{z\langle\rangle}, l\langle z_B\rangle, \llcorner l\langle z_B\rangle \lrcorner]\,\}\!\} \rightarrow$

$\{\!\{\dots, [z\langle\rangle \triangleright 0, Bob\langle News\rangle, Alice\langle News\rangle, l\langle z_B\rangle, \llcorner l\langle z_B\rangle \lrcorner]\,\}\!\}$

COMMIT:

$\{\!\{\dots, \{[\underline{z\langle\rangle \triangleright 0}, Bob\langle News\rangle, Alice\langle News\rangle, l\langle z_B\rangle, \llcorner l\langle z_B\rangle \lrcorner]\,\} \,\}\!\} \rightleftharpoons$

$\{\!\{\dots, \{[Bob\langle News\rangle, Alice\langle News\rangle, l\langle z_B\rangle, \textbf{def } z\langle\rangle \triangleright 0 \textbf{ in } 0, \llcorner l\langle z_B\rangle \lrcorner]\,\} \,\}\!\} \rightarrow$

$\{\!\{\dots, Bob\langle News\rangle, Alice\langle News\rangle, l\langle z_B\rangle\,\}\!\}$

Fig. 7. A possible computation of the Mailing list example.

through the following type system involving the set $T = \{\square_0, \square_1, \square_2\}$ of types and the following type judgements:

$\vdash P : \square_0$    The transaction primitive $[\_ : \_]$ does not appear in $P$ at all.

$\vdash P : \square_1$    Transactions may appear in $P$, but only inside definitions. $P$ does not have active transactions but it may create them after some reductions.

$\vdash P : \square_2$    $P$ has active flat transactions or may create them after some reductions.

$\vdash D : \square_0$    $D$ does not contain transactions.

$\vdash D : \square_1$    $D$ may contain flat transactions.

**Definition 3.2 (Flat (or well-typed) definitions and processes).** A definition $D$ is said *flat* or *well-typed* if $\vdash D : \square_1$ in the type system shown in Figure 8. Similarly, a process $P$ is said *flat* or *well-typed* if $\vdash P : \square_2$.

We comment on the typing rules in Figure 8. Rules (SUB-P) and (SUB-D) stand for the sub-type order $\square_0 < \square_1 < \square_2$. Clearly, the inert process 0, the emission of a message

$$
\begin{array}{ll}
\text{(Sub-P)} & \text{(Sub-D)} \\
\dfrac{\vdash P : \square_i}{\vdash P : \square_j}\ i < j & \dfrac{\vdash D : \square_0}{\vdash D : \square_1}
\end{array}
\qquad
\begin{array}{l}
\text{(Zero)} \\
\vdash 0 : \square_0
\end{array}
\qquad
\begin{array}{l}
\text{(Mess)} \\
\vdash x\langle y\rangle : \square_0
\end{array}
\qquad
\begin{array}{l}
\text{(Abort)} \\
\vdash abort : \square_0
\end{array}
$$

$$
\begin{array}{l}
\text{(Par)} \\
\dfrac{\vdash P : \square_i \qquad \vdash Q : \square_i}{\vdash P|Q : \square_i}
\end{array}
\qquad
\begin{array}{l}
\text{(Trans)} \\
\dfrac{\vdash P : \square_0 \qquad \vdash Q : \square_1}{\vdash [P : Q] : \square_2}
\end{array}
\qquad
\begin{array}{l}
\text{(Def)} \\
\dfrac{\vdash D : \square_i \qquad \vdash P : \square_j}{\vdash \mathbf{def}\ D\ \mathbf{in}\ P : \square_{max(i,j)}}
\end{array}
$$

$$
\begin{array}{l}
\text{(Conj)} \\
\dfrac{\vdash D : \square_i \qquad \vdash E : \square_i}{\vdash D \wedge E : \square_i}
\end{array}
\qquad
\begin{array}{l}
\text{(Ord-0)} \\
\dfrac{\vdash P : \square_0}{\vdash J \rhd P : \square_0}
\end{array}
\qquad
\begin{array}{l}
\text{(Ord)} \\
\dfrac{\vdash P : \square_i}{\vdash J \rhd P : \square_1}
\end{array}
\qquad
\begin{array}{l}
\text{(Merge)} \\
\dfrac{\vdash P : \square_0}{\vdash J \blacktriangleright P : \square_0}
\end{array}
$$

Fig. 8. Flat cJoin typing.

$x\langle \vec{y}\rangle$ and the constant *abort* do not contain transactions and, hence, they have type $\square_0$ (Rules Zero, Mess, Abort). The parallel composition $P|Q$ has type $\square_i$ if both $P$ and $Q$ are typed $\square_i$ (rule (Par)). Rule (Trans) prevents nesting by stating that $[P : Q]$ is typed $\square_2$ only when $P$ has no transactions (i.e., $\vdash P : \square_0$). Note that the process $P \equiv [\mathbf{def}\ a\langle\rangle \rhd [P_1 : Q_1]\ \mathbf{in}\ a\langle\rangle : Q]$ is not typable. Although $P$ is not a nested transaction, $P$ may evolve to a nested transaction as follows $P \rightarrow [\mathbf{def}\ a\langle\rangle \rhd [P_1 : Q_1]\ \mathbf{in}\ [P_1 : Q_1] : Q]$. Contrastingly, a compensation may contain transactions as part of its definition. For instance, $[x\langle y\rangle\ |\ abort : \mathbf{def}\ a\langle\rangle \rhd [P : Q]\ \mathbf{in}\ a\langle\rangle]$ has type $\square_2$ when $\vdash P : \square_0$ and $\vdash Q : \square_1$. The fact that the compensation includes a transaction as part of its definition does not compromise flatness because the compensation will run as an ordinary process after the transaction aborts. In fact, $[x\langle y\rangle\ |\ abort : \mathbf{def}\ a\langle\rangle \rhd [P : Q]\ \mathbf{in}\ a\langle\rangle] \rightarrow \mathbf{def}\ a\langle\rangle \rhd [P : Q]\ \mathbf{in}\ a\langle\rangle$, which does not introduce nesting. Rule (Def) combines the types of definitions and processes. Note that $\mathbf{def}\ D\ \mathbf{in}\ P$ is typed $\square_0$ when neither $D$ nor $P$ contain transactions, i.e., if they both have type $\square_0$. A process $\mathbf{def}\ D\ \mathbf{in}\ P$ has type $\square_1$ when transactions appear only in definitions (i.e., either in $D$ or in other local definitions occurring in $P$). Finally, $\mathbf{def}\ D\ \mathbf{in}\ P$ has type $\square_2$ when $P$ contains an active transaction.

A composed definition (i.e., a conjunction) is typed $\square_i$ only when both sub-terms have type $\square_i$ (By rule (Conj)). An ordinary definition $J \rhd P$ is well-typed when its guarded processes $P$ is well-typed. Moreover, it has type $\square_0$ when $P$ has no transactions, i.e., $\vdash P : \square_0$. Differently, a merge rule is well-typed only when $P$ has type $\square_0$ (rule (Merge)). This is required in order to avoid nesting, because the instances of $P$ will execute inside transactions.

**Example 6 (Well-typed terms).** Consider the mailing list process introduced in Example 5. Several sub-terms and their types are shown below:

$$
\begin{array}{ll}
P_1 = \mathbf{def}\ z\langle\rangle \rhd 0\ \mathbf{in}\ y\langle v,z\rangle\ |\ l\langle y\rangle & P_2 = [P_1 : l\langle y\rangle] \\
D_1 = l\langle y\rangle\ |\ tell\langle v\rangle \rhd P_2 & D_2 = l\langle y\rangle\ |\ close\langle\rangle \rhd 0 \\
\vdash P_1 : \square_0 \qquad \vdash P_2 : \square_2 \qquad \vdash D_1 : \square_1 & \vdash D_2 : \square_0 \qquad \vdash D_1 \wedge D_2 : \square_1.
\end{array}
$$

Moreover, $\vdash \mathsf{MLDef} : \square_1$ (because it does not have active transactions but it can activate them) and also $\vdash \mathsf{ML} : \square_1$.

Fig. 9. Flat solution typing.

**Example 7 (Counterexample).** Process **def** $x\langle\rangle \blacktriangleright [P : 0]$ **in** $[\textbf{def } D \textbf{ in } x\langle\rangle : 0]$ is not well-typed because it has a merge definition whose guarded process is a transaction (rule (MERGE) cannot be applied because $\nvdash x\langle\rangle \blacktriangleright [P : 0] : \Box_0$). In fact, this process can reduce in one step to **def** $x\langle\rangle \blacktriangleright [P : 0]$ **in** $[\textbf{def } D \textbf{ in } [P : 0] : 0]$ when $x \notin dn(D)$, which has nested transactions.

### 3.3. *Properties of flat* cJoin *typing*

This section summarises the main properties of our type system, namely, Join processes have type $\Box_0$ (Proposition 3.3) and subject reduction holds for $\Box_0$ (Lemma 3.8) and $\Box_2$ (Lemma 3.9).

**Proposition 3.3 (Join processes have type $\Box_0$).** Let $P$ be a Join process, then $\vdash P : \Box_0$.

*Proof.* The proof follows by induction on the structure of $P$.

— $P \equiv 0$ and $P \equiv x\langle\vec{y}\rangle$: the proof follows by using either rule (ZERO) or (MESS).
— $P \equiv \textbf{def } D \textbf{ in } P'$ with $D = \wedge_i J_i \triangleright P_i$. By inductive hypothesis, $\vdash P' : \Box_0$ and $\vdash P_i : \Box_0$ for all $i$. By using rule (ORD-0), we conclude that $\vdash J_i \triangleright P_i : \Box_0$ for all $i$. By repeatedly using rule (Conj) we conclude $\vdash D : \Box_0$. Proof is completed by applying rule (DEF).
— $P \equiv P_1|P_2$: the proof follows by inductive hypothesis and rule (PAR). $\qquad\square$

In order to prove subject reduction we need some technical preliminaries. In particular, we extend the typing from processes to solutions.

**Definition 3.4 (Type of a solution).** The type $\tau$ of a solution $S$, noted as $S : \tau$, is defined by rules in Figure 9. Moreover, $S$ is flat iff $S : \Box_2$.

We start by proving that Definition 3.4 is consistent w.r.t. structural congruence of solutions, i.e. all types are preserved by $\alpha$-conversion and heating/cooling; and that the type of a solution reflects on the type of its molecules.

**Proposition 3.5.** Let $\sigma$ be a renaming substitution. If $\vdash P : \tau$ then $\vdash P\sigma : \tau$.

*Proof.* Immediate by the fact that typing does not take into account names, but just the structure of terms, which cannot be changed by renaming substitutions. $\qquad\square$

**Lemma 3.6.** Let $S : \Box_j$. If $\{|S|\} \rightleftharpoons \{|S'|\}$ then $S' : \Box_j$.

*Proof.* By straightforward case analysis on the applied cooling/heating rule. When the applied rule is (STR-DEF), then Proposition 3.5 is used. $\qquad\square$

**Corollary 3.7.** Let $\{[S]\} \rightleftharpoons^* \{[\otimes_i m_i]\}$. Then $S : \square_j$ iff $\forall i. m_i : \square_i$ and $i \leqslant j$.

We are now ready to prove subject reduction for $\square_0$.

**Lemma 3.8 (Subject Reduction for $\square_0$).** Let $P : \square_0$. If $P \rightarrow^* P'$ then $P' : \square_0$.

*Proof.* The proof follows by induction on the length of the derivation.

— **Base case**: $P' \equiv P$ (i.e., $\{[P]\} \rightleftharpoons^* \{[P']\}$). The proof follows by Corollary 3.7.
— **Inductive Step**: suppose $P \rightarrow P'' \rightarrow^n P'$ with $n \geqslant 0$. The proof follows by case analysis on the first applied rule and inductive hypothesis. Note that $P \rightarrow P''$ implies $\{[P]\} \rightleftharpoons^* \{[S]\} \rightarrow \{[S']\} \rightleftharpoons^* \{[P'']\}$. By CHAM semantics we know that $S \equiv \otimes_i m_i$. Since $P : \square_0$, then $m_i : \square_0$ for all $i$ by Corollary 3.7. Hence, the only possible rule that can be applied is (RED), because any other rule requires at least a molecule composed by a membrane, which cannot be typed $\square_0$. Consequently, $S \equiv J \triangleright Q, J\sigma, S''$ and $S' \equiv J \triangleright Q, Q\sigma, S''$, where $S'' : \square_0$. As $Q : \square_0$, by Proposition 3.5, $Q\sigma : \square_0$. Hence, $S' : \square_0$ and therefore $P'' : \square_0$ (by Corollary 3.7). The proof follows by applying inductive hypothesis on $P'' \rightarrow^n P'$.

□

The following result ensures that $\square_2$ is preserved by reductions.

**Theorem 3.9 (Subject reduction for $\square_2$).** Let $P : \square_2$. If $P \rightarrow^* P'$ then $P' : \square_2$.

*Proof.* The proof follows by induction on the length of the derivation.

— **Base case**: $P' \equiv P$ (i.e., $\{[P]\} \rightleftharpoons^* \{[P']\}$). The proof follows by Corollary 3.7.
— **Inductive Step**: suppose $P \rightarrow P'' \rightarrow^n P'$ with $n \geqslant 0$. The proof follows by case analysis on the first applied rule and inductive hypothesis. Note that $P \rightarrow P''$ implies $\{[P]\} \rightleftharpoons^* \{[S]\} \rightarrow \{[S']\} \rightleftharpoons^* \{[P'']\}$. By CHAM semantics we know that $S \equiv \otimes_i m_i$. Since, $P : \square_0$, $m_i : \square_i$ where $i \leqslant 2$ for all $i$ by Corollary 3.7. Then, there are four cases:

  – Rule (RED): when the reduction occurs at top-level, i.e., $S \equiv J \triangleright Q, J\sigma, S''$, $S' \equiv J \triangleright Q, Q\sigma, S''$, and $S'' : \square_2$, the proof is similar to Lemma 3.8. The other possibility is when the reduction occurs inside a transaction, e.g., $S \equiv \{[S_1]\}, S''$ and $S'' \equiv \{[S_1']\}, S''$, where $\{[S_1]\} \rightarrow \{[S_1']\}$ by rule (RED) and $S'' : \square_2$. Note that $\{[S_1]\} : \square_2$, and therefore $S_1 : \square_0$. By Lemma 3.8, $S_1' : \square_0$ and hence $S' : \square_2$.

(The cases below occur at top-level, because negotiations cannot be nested in $P$.)

  – Rule (COMMIT): $S \equiv \{[M | \mathbf{def}\ D\ \mathbf{in}\ 0, {}_{\llcorner}Q_{\lrcorner}]\}, S''$, and $S' \equiv M, S''$, with $S'' : \square_2$ (by Corollary 3.7). As $M$ is the parallel composition of messages, it can be typed $\square_0$ and therefore $S' : \square_2$.

  – Rule (ABORT): $S \equiv \{[abort | P', {}_{\llcorner}Q_{\lrcorner}]\}, S''$ and $S' \equiv Q, S''$, with $S'' : \square_2$ (by Corollary 3.7). As ${}_{\llcorner}Q_{\lrcorner} : \square_0$, it must be $Q : \square_1$ and therefore $S' : \square_2$.

  – Rule (MERGE): $S \equiv J_1 | \ldots | J_n \blacktriangleright R, \otimes_i \{[J_i \sigma, S_i, {}_{\llcorner}Q_i{}_{\lrcorner}]\}, S''$ and $S' \equiv J_1 | \ldots | J_n \blacktriangleright R, \{[\otimes_i S_i, R\sigma, {}_{\llcorner}Q_1 | \ldots | Q_n{}_{\lrcorner}]\}, S''$, with $S'' : \square_2$ (by Corollary 3.7). Since $R : \square_0$ and for all $i$ $S_i : \square_0$ and $Q_i : \square_1$, we have $S' : \square_2$.

□

**Remark 3.10.** Subject reduction does not hold for $\square_1$. Consider $P = \mathbf{def}\ x\langle\rangle \triangleright [Q : Q']\ \mathbf{in}\ x\langle\rangle$, where $\vdash Q : \square_0$ and $\vdash Q' : \square_1$. Although $\vdash P : \square_1$, $P$ reduces to $P' = \mathbf{def}\ x\langle\rangle \triangleright [Q : Q']\ \mathbf{in}\ [Q : Q']$, which can be typed $\square_2$ but not $\square_1$.

Theorem 3.9 ensures that flat processes form a sub-calculus since reductions do not generate nesting.

**Definition 3.11 (Flat cJoin).** Flat cJoin is the sub-calculus of all flat processes.

## 4. Programming common transactional patterns in cJoin

In order to illustrate the transactional aspects of cJoin, we show how to code some common interaction patterns. We keep and enrich the hotel booking scenario as a running example.

### 4.1. *Multi-step transactions: trip booking*

Let us assume now that the user is making plans for a trip and wants to make reservations for both a flight and hotel accommodation. Such scenarios are usually modelled by splitting the whole activity as a sequence of two independent transactions: the client executes a transaction for booking a flight first, and when it commits, a new transaction for making hotel reservation is started. If the last transaction aborts, then the first one is compensated for by cancelling the already committed flight reservation. This kind of composition is usually referred to as compensable flow composition, and it is tailored to model long running transactions in an orchestration context. Several proposals have appeared in the literature for describing transactional flow compositions (Bruni *et al.* 2005; Butler and Ferreira 2004; Butler *et al.* 2005b). Essentially, they describe the trip booking problem as follows:

$$\mathrm{TripBooking} \equiv \{\mathrm{FlightBooking} \div \mathrm{FlightCancelation}; \mathrm{BookHotel} \div 0\}.$$

The entire activity TripBooking is delimited by the long running transaction scope $\{\_\}$. A long running transaction is divided by ';' into sequential steps. Process FlightBooking in the first step allows the client to book a hotel accommodation. If it commits, then the compensation FlightCancelation is installed and the next step $\mathrm{BookHotel} \div 0$ is executed. The compensation is installed when a step commits and it is used only when one of the following steps fails. For instance, if BookHotel fails during its execution, then the previously installed compensation FlightCancelation is executed.

Although cJoin does not offer a built-in mechanism for these kinds of transactions, they can be coded into cJoin. Let us consider a simple language defined as follows:

$$
\begin{array}{rcl}
L & ::= & \{S\} \\
S & ::= & P \mid P \div Q; S
\end{array}
$$

where, $P$ and $Q$ are cJoin processes. A possible encoding $[\![\_]\!]$ is given below:

$$
\begin{array}{lll}
[\![\{S\}]\!] & \equiv & [\![S]\!]_0 \\
[\![P]\!]_c & \equiv & [P : c] \\
[\![P \div Q; S]\!]_c & \equiv & \mathbf{def}\ comp\langle\rangle \triangleright [c \mid Q : 0] \\
& & \quad \wedge\ cont\langle\rangle \triangleright [\![S]\!]_{comp\langle\rangle} \\
& & \mathbf{in}\ [cont\langle\rangle \mid P : c] \\
& & \qquad\qquad provided\ \{cont, comp\} \cap fn(P|Q|S|c) = \varnothing.
\end{array}
$$

The most interesting rule is the last one. A sequence $P \div Q; S$ is encoded w.r.t. a context $c$ that indicates the compensation installed by the previous execution. Then, the sequence $P \div Q; S$ corresponds to a process that activates the transaction $[cont\langle\rangle \mid P : c]$, i.e., it attempts to executes $P$ until completion (i.e., consuming all messages to its local ports). If $P$ finishes (i.e., there is no pending message in local ports) and the transaction commits, then the message $cont\langle\rangle$ is released. This message will activate the execution of the remaining part of the sequence. If $P$ aborts, then the previously installed compensation $c$ is activated. The context for encoding the remaining part $S$ of the sequence is the message $comp\langle\rangle$, corresponding to the updated compensation $[c \mid Q : 0]$. Hence, if activated, such compensation will first attempt to complete the execution of $Q$. If this is the case, then the transaction will commit by releasing $c$, which will then activate the previously installed compensations.

Then, the cJoin code for planning a trip as a multi-step transaction is in Figure 10. We model the hotel booking service H as in Figure 6, while the airline service A is analogous to H (the only difference is that A starts by publishing on the merge channel *AirlineSrv* instead of *HotelSrv*). It is worth remarking that the homonymous (private) ports *request* and *accept* defined by both A and H are different ports. The renaming mechanism inherited from the Join calculus ensures that such names are fresh, and hence there are no clashes when transactions are merged. In addition, we add one merging rule for allowing the interaction among the client and the airline. For the sake of the simplicity, we do not include dynamic creation of sessions, but the presentation can be straightforwardly extended to consider them.

The system starts its execution with the client interacting with the airline component analogously to hotel conversation described in Example 4. If the interaction aborts, then the client finishes its execution (note that transaction compensation is set to 0). If it commits, the message $cont\langle\rangle$ is released, which enables the second reaction rule of C′. Firing this rule will dynamically create a new transaction for interacting with the hotel booking service (the interaction is as in Example 4). If this newly created transaction finishes successfully, then the whole long running transaction ends by committing. If the last transaction aborts, then the compensation $comp\langle\rangle|P_2$ is released. Note that $comp\langle\rangle$ can fire the first reaction of client component, which will execute FlightCancelation, i.e., the compensation of the previously committed step.

$$
\begin{aligned}
\mathsf{TB} \ &\equiv \ \mathbf{def} \ HotelSrv\langle r\rangle \mid HotelReq\langle d,\kappa\rangle \ \blacktriangleright \ r\langle d,\kappa\rangle \\
&\qquad \wedge \ AirlineSrv\langle r\rangle \mid AirlineReq\langle d,\kappa\rangle \ \blacktriangleright \ r\langle d,\kappa\rangle \\
&\qquad \mathbf{in} \ \mathsf{A} \mid \mathsf{H} \mid \mathsf{C}' \\[6pt]
\mathsf{A} \ &\equiv \ [\mathbf{def} \ request\langle details,\kappa\rangle \ \triangleright \ \kappa\langle priceFlight,accept\rangle + abort \\
&\qquad \wedge \ accept\langle cc\rangle \ \triangleright \ 0 \\
&\qquad \mathbf{in} \ AirlineSrv\langle request\rangle : P_1] \\[6pt]
\mathsf{H} \ &\equiv \ [\mathbf{def} \ request\langle details,\kappa\rangle \ \triangleright \ \kappa\langle priceRoom,accept\rangle + abort \\
&\qquad \wedge \ accept\langle cc\rangle \ \triangleright \ 0 \\
&\qquad \mathbf{in} \ HotelSrv\langle request\rangle : P_2] \\[6pt]
\mathsf{C}' \ &\equiv \ \mathbf{def} \ comp\langle\rangle \ \triangleright \ [\mathtt{FlightCancellation} : 0] \\
&\qquad \wedge \ cont\langle\rangle \ \triangleright \ [\mathbf{def} \ offer\langle rate,\kappa\rangle \ \triangleright \ \kappa\langle card\rangle + abort \\
&\qquad\qquad\qquad\qquad\quad \mathbf{in} \ HotelReq\langle data,offer\rangle : comp\langle\rangle] \\
&\qquad \mathbf{in} \ [cont\langle\rangle \\
&\qquad\quad \mid \mathbf{def} \ offer\langle rate,\kappa\rangle \ \triangleright \ \kappa\langle card\rangle + abort \\
&\qquad\qquad \mathbf{in} \ AirlineReq\langle data,offer\rangle : 0]
\end{aligned}
$$

Fig. 10. Trip booking as a multi-step transaction.

### 4.2. *Multi-way transactions: trip booking revisited*

The main drawback of planning a trip as in the previous section is that compensations are usually not for free. Clearly, cancelling a flight reservation usually requires the client to pay cancelation fees. By taking advantage of cJoin transactional primitives, we can model the trip booking example as a multi-way transaction, i.e., a transaction that retains several entry and exit points, in which parties are not necessarily aware of the remaining parties in the transaction. Clearly, the hotel and airline booking services should not be necessarily aware of the combined activity to be carried on by the client, and hence they remain unchanged. The cJoin process for the whole system is shown in Figure 11.

The client is modelled by a transaction that initially sends two messages to different merge channels, one allows the merge with the airline service while the other joins it to the hotel service. Once merged, the parties interact by following the conversation pattern described in Example 4. The main difference is that the abortion of one of the parties after both merge actions take place implies the abortion of the interaction, and hence all three parties abort and run their own compensations. Otherwise, the whole interaction commits when all three parties successfully finish their transactional processes.

As far as the definition of the choreography is concerned, we note that (i) booking services are independent of the behaviour of the client, even though this may induce the transactional scope to be extended to third, unknown parties; (ii) the local description of transactional interfaces ensures transactional properties to several different global interactions; (iii) no coordinator is needed for describing transactional, multiparty choreographies.

There is still one main drawback in the code given above: client description mixes two independent flows of interaction corresponding to two different roles played by the component either as a hotel client or as an airline client. The only binding among these

$$\begin{aligned}
\mathsf{TB} \quad \equiv \quad & \textbf{def} \quad HotelSrv\langle r\rangle \mid HotelReq\langle d,\kappa\rangle \ \blacktriangleright\ r\langle d,\kappa\rangle \\
& \wedge \ AirlineSrv\langle r\rangle \mid AirlineReq\langle d,\kappa\rangle \ \blacktriangleright\ r\langle d,\kappa\rangle \\
& \textbf{in} \ \mathsf{A} \mid \mathsf{H} \mid \mathsf{C}'
\end{aligned}$$

$$\begin{aligned}
\mathsf{A} \quad \equiv \quad & [\textbf{def} \ request\langle details,\kappa\rangle \ \triangleright\ \kappa\langle priceFlight, accept\rangle + abort \\
& \wedge \ accept\langle cc\rangle \ \triangleright\ 0 \\
& \textbf{in} \ AirlineSrv\langle request\rangle : P_1]
\end{aligned}$$

$$\begin{aligned}
\mathsf{H} \quad \equiv \quad & [\textbf{def} \ request\langle details,\kappa\rangle \ \triangleright\ \kappa\langle priceRoom, accept\rangle + abort \\
& \wedge \ accept\langle cc\rangle \ \triangleright\ 0 \\
& \textbf{in} \ HotelSrv\langle request\rangle : P_2]
\end{aligned}$$

$$\begin{aligned}
\mathsf{C}' \quad \equiv \quad & [\textbf{def} \ hotelOffer\langle rate,\kappa\rangle \ \triangleright\ \kappa\langle card\rangle + abort \\
& \wedge \ airOffer\langle rate,\kappa\rangle \ \triangleright\ \kappa\langle card\rangle + abort \\
& \textbf{in} \ \ HotelReq\langle dataRoom, hotelOffer\rangle \\
& \quad \mid AirlineReq\langle dataFlight, airOffer\rangle : Q]
\end{aligned}$$

Fig. 11. Trip booking.

two flows of execution is the fact that both interactions should finish either successfully or with abortion. A more appealing modular definition for client behaviour is below:

$$\begin{aligned}
\mathsf{C_m} \quad \equiv \quad \textbf{def} \ \ & RoomFound\langle\omega\rangle \mid FlightFound\langle\omega\rangle \ \blacktriangleright\ 0 \\
\textbf{in} \ \ & [\textbf{def} \ hotelOffer\langle rate,\kappa\rangle \ \triangleright\ \kappa\langle card\rangle \mid RoomFound\langle hotelOffer\rangle + abort \\
& \quad \textbf{in} \ \ HotelReq\langle dataRoom, hotelOffer\rangle : Q_1] \\
& \mid [\textbf{def} \ airOffer\langle rate,\kappa\rangle \ \triangleright\ \kappa\langle card\rangle \mid FlightFound\langle airOffer\rangle + abort \\
& \quad \textbf{in} \ AirlineReq\langle dataFlight, airOffer\rangle : Q_2].
\end{aligned}$$

Now, a client initiates two different transactions to interact independently with the hotel and the airline. The new merge rule allows such transactions to be joined when both flows of interaction terminate. In fact, the corresponding messages to the merge names (i.e., *RoomFound*⟨*hotelOffer*⟩ and *FlightFound*⟨*airOffer*⟩) are generated when the client transactions accept the offers proposed by the corresponding booking services. Let us consider the monolithic code again (Figure 11). For instance, if the hotel offers a convenient rate but the airline company does not, then the whole transaction is aborted. Hence, the client should start from the scratch by booking again a room and trying with a different airline company. In the modular version, it would be enough to start a new transaction for finding a flight, while the booked room will be still a valid reservation. This feature may be interesting when client also may choose one of several available booking services (as discussed in the next sections). For example, the client can try a new airline booking (after an abort) while keeping the room booking. Note that this example also shows how a typical nested transaction pattern is smoothly modelled as a flat cJoin process.

A final remark is that the above decomposition of flows works only when both interactions eventually finish by proposing the merge, otherwise one transaction may remain blocked for ever. When this is not ensured, the interaction that fails to find a suitable reservation should notify the other transaction to abort (this situation can be handled analogously to abortions of parallel activities explained in Section 4.3).

### 4.3. *Speculative computation*

This section illustrates how to program a cJoin process that commits at most one of several concurrent transactions. This programming pattern is called speculative execution or a posteriori choice, and it is studied by several composition languages (Bruni *et al.* 2005; Laneve and Zavattaro 2005; Butler *et al.* 2005b). Speculative computation is very much related to goal-oriented formalisms, like don't know non-determinism in concurrent logic programming, where several alternatives must be explored before one can be selected. For example, when guarded Horn clauses are considered, the selection of the rule to be applied next for reducing a goal is subordinated to successful evaluation of the guard. Once a guarded clause is applied, all the other alternatives are pruned out and that intermediate goal reduction is committed (never to be undone). As guard evaluation can possibly trigger complex computations on its own and require further clause selections giving rise to speculative computations, whose abort corresponds to the non-satisfiability of the clause. One important difference though is that the linguistic abstraction we are after is not a basic search mechanism, deemed to fail in most cases, but rather a strategy language to increase the possibility of success and to handle search failures when they happen.

Although this pattern is not a built-in operator of cJoin, it can be coded by using merge definitions. For simplicity, we will consider just two concurrent transactions, but the mechanism can be easily extended to larger sets.

We assume a client trying to book a room from one of two alternative hotels $H_1$ and $H_2$ (in no particular order), but wishing to make a reservation in only one of them. The system is described below:

$$
\begin{aligned}
\mathsf{HB} \;\equiv\; &\textbf{def}\;\; HotelSrv_1\langle r \rangle \mid HotelReq_1\langle d, \kappa \rangle \;\blacktriangleright\; r\langle d, \kappa \rangle && (1)\\
&\quad\quad HotelSrv_2\langle r \rangle \mid HotelReq_2\langle d, \kappa \rangle \;\blacktriangleright\; r\langle d, \kappa \rangle && (2)\\
&\textbf{in}\;\; \mathsf{H}_1 \mid \mathsf{H}_2 \mid \mathsf{C} && (3)
\end{aligned}
$$

$$
\begin{aligned}
\mathsf{C} \;\equiv\; &\textbf{def}\;\; alt\langle id \rangle \mid booked_1\langle w \rangle \;\blacktriangleright\; cancelling_2\langle\rangle && (4)\\
&\;\wedge\; alt\langle id \rangle \mid booked_2\langle w \rangle \;\blacktriangleright\; cancelling_1\langle\rangle && (5)\\
&\;\wedge\; cancelling_1\langle\rangle \rhd [\textbf{def}\; w\langle\rangle \rhd 0 \;\textbf{in}\; aborting_1\langle w \rangle : 0] && (6)\\
&\;\wedge\; cancelling_2\langle\rangle \rhd [\textbf{def}\; w\langle\rangle \rhd 0 \;\textbf{in}\; aborting_2\langle w \rangle : 0] && (7)\\
&\;\wedge\; booked_1\langle i \rangle \mid aborting_1\langle j \rangle \;\blacktriangleright\; abort && (8)\\
&\;\wedge\; booked_2\langle i \rangle \mid aborting_2\langle j \rangle \;\blacktriangleright\; abort && (9)\\
&\;\wedge\; a\langle\rangle \rhd [\textbf{def}\; w\langle\rangle \rhd 0 \;\textbf{in}\; alt\langle w \rangle : a\langle\rangle] && (10)\\
&\textbf{in}\;\; [\textbf{def}\; offer\langle rate, \kappa \rangle \;\rhd\; \kappa\langle card \rangle \mid booked_1\langle offer \rangle && (11)\\
&\quad\quad\quad\quad\quad\quad\quad\quad\quad + abort && (12)\\
&\quad\quad \textbf{in}\; HotelReq_1\langle data, offer \rangle \quad : Q_1] && (13)\\
&\;\mid\; [\textbf{def}\; offer\langle rate, \kappa \rangle \;\rhd\; \kappa\langle card \rangle \mid booked_2\langle offer \rangle && (14)\\
&\quad\quad\quad\quad\quad\quad\quad\quad\quad + abort && (15)\\
&\quad\quad \textbf{in}\; HotelReq_2\langle data, offer \rangle \quad : Q_2] && (16)\\
&\;\mid\; a\langle\rangle. && (17)
\end{aligned}
$$

Lines (1)–(3) define the system, which is composed by hotel services $H_1$ and $H_2$ defined as in previous examples, and a client component C. The client C consists in the concurrent execution of two transactions (lines (11)–(13) and (14)–(16)) for dealing with each booking service. They behave as in previous cases, but they finish the conversation by generating the merge messages $booked_i\langle offer \rangle$ (lines (11) and (12)). Differently from multi-way transactions (Section 11), those messages are used not for joining both transactions but for encoding a kind of transactional internal choice. Let us suppose that both client transactions complete successfully; then the state of the client can be seen as follows:

$$
\begin{aligned}
&\mathbf{def}\ \dots && (4-10)\\
&\mathbf{in}\quad [\mathbf{def}\ \dots && (11-12)\\
&\qquad \mathbf{in}\ booked_1\langle offer \rangle \quad : Q_1] && (13')\\
&\quad \mid [\mathbf{def}\ \dots && (14-15)\\
&\qquad \mathbf{in}\ booked_2\langle offer \rangle \quad : Q_2] && (16')\\
&\quad \mid [\mathbf{def}\ w\langle \rangle \rhd 0\ \mathbf{in}\ alt\langle w \rangle : a\langle \rangle] && (17')
\end{aligned}
$$

where definitions rules have been omitted since they remain unchanged. Line (17′) has been obtained by firing the reaction rule in line (10) with message $a\langle \rangle$ (line (17)). No transaction can commit in this state because messages sent to merge ports carry on local names (*offer* and $w$) – this is the standard way to force a transaction to join before committing. At this point, the merge rules in lines (4) and (5) are enabled. Assuming the second one fires, the client state reduces to

$$
\begin{aligned}
&\mathbf{def}\ \dots && (4-10)\\
&\mathbf{in}\quad [\mathbf{def}\ \dots && (11-12)\\
&\qquad \mathbf{in}\ booked_1\langle offer \rangle \quad : Q_1] && (13')\\
&\quad \mid [\mathbf{def}\ \dots && (14-15)\\
&\qquad \wedge\ w\langle \rangle \rhd 0 &&\\
&\qquad \mathbf{in}\ cancelling_1\langle \rangle : a\langle \rangle | Q_2]. && (16')
\end{aligned}
$$

Assuming the hotel component also commits, then the second transaction commits by releasing the message $cancelling_1\langle \rangle$, which enables the reaction rule in line (6). After firing this rule, client state can be described as follows:

$$
\begin{aligned}
&\mathbf{def}\ \dots && (4-10)\\
&\mathbf{in}\quad [\mathbf{def}\ \dots && (11-12)\\
&\qquad \mathbf{in}\ booked_1\langle offer \rangle \quad : Q_1] && (13')\\
&\quad \mid [\mathbf{def}\ w\langle \rangle \rhd 0\ \mathbf{in}\ aborting_1\langle w \rangle : 0].
\end{aligned}
$$

Above state enables the merge rule in line (8). Client component moves to the following state when such rule is fired.

$$
\begin{aligned}
&\textbf{def} \ldots &&(4-10)\\
&\quad \textbf{in} \quad [\textbf{def} \ldots &&(11-12)\\
&\qquad\qquad \wedge\, w\langle\rangle \rhd 0\\
&\qquad\quad \textbf{in}\ abort \quad : Q_1].
\end{aligned}
$$

Finally, the remaining transaction aborts and the compensation $Q_1$ is released. Hence, the transaction with $\mathsf{H}_2$ has been committed while the conversation with $\mathsf{H}_1$ has been aborted. The cases in which one transaction commits and the other aborts follow immediately since rules defined in lines (8) and (9) are never used, and therefore, the auxiliary transaction generated by message *cancelling*$_i$ remains blocked.

We comment on reaction rule in line (10), which generates an auxiliary transaction for selecting the committing interaction. Note that its compensation is set to $a\langle\rangle$. Although, the scenario described so far assumes that the hotel component does not generate *abort* after receiving client confirmation, this may not be the case in the most general setting. Hence, if the selected transaction aborts after being chosen for committing, then the compensation $a\langle\rangle$ is released in order to allow the remaining alternative to be eligible again.

## 5. Language implementation

This section addresses the problem of implementing the transactional primitives provided by cJoin. For simplicity, we only consider flat cJoin, i.e., the sub-calculus of transactions without nesting introduced in Definition 3.11. In particular, we show how flat cJoin can be encoded into Join. For the sake of the simplicity, we omit many technical details in this presentation and provide just a sketch of the translation and the main results. The formal definition of the encoding and its completeness and correctness results can be found in (Melgratti 2005).

Intuitively, transactional processes are implemented in Join by making explicit a commit protocol used by several parties to reach an agreement. In particular, we rely on the *distributed two phase commit* (D2PC) of (Bruni *et al.* 2002), because it is appropriate for handling situations in which parties are not necessary aware of the whole set of participants involved in the transaction. Nevertheless, the encoding is parametric w.r.t. to the selected commit protocol, hence, the translation can be adapted to make participants conform to other proposals such as the standards WS-ATOMIC TRANSACTION or WS-BUSINESS ACTIVITY. In what follows we call a *coordinator* any party performing the selected commit protocol. We will denote coordinators by Coor. The formal definition for the Join processes Coor used in our translation can be found in (Melgratti 2005). The code is also reported in the appendix for the interested reader, but we find unnecessary here to illustrate its code in detail, because the D2PC is not the focus of this paper. Here, let us just assume that a coordinator offers (a fresh instance of) the following ports to communicate with:

| Name | Parameters | Stands for |
|------|------------|------------|
| *cmp* | *m* | to set the compensation *m* to be delivered on abort |
| *cmt* | $\kappa$ | to start the protocol by voting commit. $\kappa$ is the continuation to be released on commit |
| *abt* | | to start the protocol by voting abort |
| *join* | *coord* | to join the coordinator *coord* to the same transaction |

The key strategy for encoding a transaction into a Join process is to assign a coordinator to any execution thread of the transaction, i.e., any message sent to a transactional or merge port is monitored by one coordinator. This is achieved by making any such message carry on the ports that allows ones to interact with its coordinator. For instance, consider the following cJoin transaction:

$$T \equiv [\ \mathbf{def}\ x\langle\rangle \triangleright P\ \mathbf{in}\ x\langle\rangle\ :Q].$$

The corresponding Join term, denoted $[\![T]\!]$, is defined as follows:

$$[\![T]\!] \equiv \mathbf{def}\ \mathsf{Coor}$$
$$\wedge\ x\langle c, a, j\rangle \triangleright [\![P]\!]_{c,a,j}$$
$$\wedge\ undo\langle\rangle \triangleright [\![Q]\!]$$
$$\mathbf{in}\ \ x\langle cmt, abt, join\rangle$$
$$|\ cmp\langle undo\rangle.$$

The transaction $T$ is encoded as the process $[\![T]\!]$, which introduces one fresh coordinator Coor to monitor the unique execution thread of $T$ (i.e., message $x\langle\rangle$ in this particular case). As mentioned before, Coor defines the fresh ports *cmt*, *abt*, *join* and *cmp*. The message $x\langle\rangle$ in $T$, which is monitored by Coor in the translation, is encoded as the message $x\langle cmt, abt, join\rangle$, which carries on the names needed for interacting with Coor (the usage of these names will be illustrated below). Note that the original definition of $x$ (rule $x\langle\rangle \triangleright P$) needs to be amended to take into account the three new parameters, i.e., the original definition is mapped to $x\langle c, a, j\rangle \triangleright [\![P]\!]_{c,a,j}$. We highlight that the original guarded process $P$ is translated as $[\![P]\!]_{c,a,j}$, where $[\![P]\!]_{c,a,j}$ denotes the encoding of the process $P$ that is being monitored by a coordinator identified by the ports $c, a, j$. Finally, we introduce the new local port *undo*, which is used as the guard of the (encoded form of the) compensation $Q$. Note that this name is used for setting up the compensation of Coor (message $cmp\langle undo\rangle$). We assume that Coor will deliver a message to its settled compensation when the execution of the commit protocol finishes with abort. Consequently, rule $undo\langle\rangle \triangleright [\![Q]\!]$ will be enabled only when the commit protocol aborts. When fired, the encoded form of the original compensation $Q$ is activated.

We now analyse how a monitored processes $P$ is translated into $[\![P]\!]_{c,a,j}$. We start by considering the encoding of the following five different forms of monitored processes.

— $\mathbf{P} \equiv \mathbf{0}$. In this case the monitored process $P$ has finished its execution successfully. Consequently, this thread can request its coordinator to initiate the commit protocol

by voting commit. Hence, the encoding of 0 is defined as follows:

$$[\![0]\!]_{c,a,j} = c\langle\rangle.$$

The encoded form of 0 is a message sent to the commit port of the monitor of the thread. This message causes the coordinator to start the commit protocol by proposing commit. If all involved parties in the transaction commit, then the transaction finishes successfully. The coordinator monitoring $P$ will finish silently, since no continuation is being set when message $c\langle\rangle$ is sent (i.e., $c\langle\rangle$ does not carry on any value).

— $P = abort$. In this case, the thread being monitored reaches the abort condition. Consequently, it informs its coordinator that the whole transaction must abort by sending a message to the corresponding abort port. Hence, *abort* is encoded as follows:

$$[\![abort]\!]_{c,a,j} = a\langle\rangle.$$

In this case, the commit protocol will finish with abort because there is at least one coordinator voting for abort. As a consequence, the coordinators in the transaction will release their settled compensations. In fact, the message $cmp\langle undo\rangle$ will be consumed to issue $undo\langle\rangle$ and therefore trigger $[\![Q]\!]$.

— **P is a message sent to a transactional or merge port**. The encoding of a monitored message sent to a transactional or a merge port is just obtained by extending the parameters of the original message with the ports corresponding to the monitor of the thread (as done for the initial thread of a transaction). Therefore, the encoding of a transactional message is defined below:

$$[\![x\langle\vec{v}\rangle]\!]_{c,a,j} = x\langle c, a, j, \vec{v}\rangle.$$

— **P is a message sent to a non transactional port**. As an example, consider the transaction $T \equiv [\ \mathbf{def}\ x\langle\rangle \triangleright P\ \mathbf{in}\ x\langle\rangle\ |\ R : Q]$ with $P = z\langle\rangle$. A possible reduction for $T$ is $T \rightarrow^* [\ \mathbf{def}\ x\langle\rangle \triangleright P\ \mathbf{in}\ z\langle\rangle\ |\ R : Q]$. Note that the message $z\langle\rangle$ produced inside of the transaction will be delivered to its recipient only when the transaction commits. Consequently, a rule like $x\langle\rangle \triangleright z\langle\rangle$ above should be interpreted as a transactional thread that is finishing its execution and expects to deliver $z\langle\rangle$ if the transaction finally commits. Hence, its encoding is defined as follows:

$$[\![z\langle\rangle]\!]_{c,a,j} = c\langle z\rangle.$$

The main difference with the encoding of 0 is that the commit message sent to the coordinator sets $z$ as the continuation. The coordinator will start the commit protocol by voting commit after receiving this message. If the commit protocol finally terminates with commit, then this coordinator will release the message $z\langle\rangle$.

— **P has two execution threads**. For instance, $P = P_1|P_2$. In this case, the translation needs to dynamically generate a new coordinator because any thread needs to be monitored by one coordinator. The encoding is defined as follows:

$$[\![P_1|P_2]\!]_{c,a,j} = \mathbf{def}\ \mathsf{Coor}\ \mathbf{in}\ [\![P_1]\!]_{cmt,abt,join}\ |\ [\![P_2]\!]_{c,a,j}\ |\ j\langle cmt, abt\rangle\ |\ join\langle c, a\rangle.$$

Note that we generate a new coordinator $\mathsf{Coor}$ that provides the definition for the fresh ports *cmt*, *abt* and *join*. Then, $P_1$ will be monitored by the newly defined

coordinator (i.e., $\llbracket P_1 \rrbracket_{cmt,abt,join}$) while $P_2$ will be monitored by the coordinator already assigned to the whole process $P_1|P_2$ (i.e., $\llbracket P_2 \rrbracket_{c,a,j}$). Messages $j\langle cmt, abt \rangle$ and $join\langle c, a \rangle$ make coordinators aware of each other: $j\langle cmt, abt \rangle$ joins the new coordinator as a participant of the transaction being monitored by the coordinator identified by $c, a, j$, while $join\langle c, a \rangle$ works in the other way round.

We now focus on the encoding of firing rules. The encoding of reduction rules outside of a transaction is simply the application of the encoding to the guarded process, i.e.,

$$\llbracket J \triangleright P \rrbracket = J \triangleright \llbracket P \rrbracket.$$

The translation treats similarly transactional and merge rules and we just consider two different shapes for such kind of rules:

— **Single-message join pattern.** Such rules are either like $x\langle \vec{v} \rangle \triangleright P$ or $x\langle \vec{v} \rangle \blacktriangleright P$. As mentioned before, rule $x\langle \vec{v} \rangle \triangleright P$ is encoded by adding to the port $x$ three new parameters for identifying the coordinator associated with the thread. Moreover, that coordinator will monitor also the execution of the guarded process $P$. The translation is defined as follows:

$$\llbracket x\langle \vec{v} \rangle \triangleright P \rrbracket = x\langle c, a, j, \vec{v} \rangle \triangleright \llbracket P \rrbracket_{c,a,j}.$$

Analogously, the encoding of a merge rule is as follows:

$$\llbracket x\langle \vec{v} \rangle \blacktriangleright P \rrbracket = x\langle c, a, j, \vec{v} \rangle \triangleright \llbracket P \rrbracket_{c,a,j}.$$

— **Two-message join pattern.** Reactions contain synchronisations, e.g., $x\langle \rangle | y\langle \rangle \blacktriangleright P$. This rule is encoded as follows:

$$x\langle c_1, a_1, j_1 \rangle \mid y\langle c_2, a_2, j_2 \rangle \triangleright \llbracket P \rrbracket_{c_1,a_1,j_1} \mid j_1\langle c_2, a_2 \rangle \mid j_2\langle c_1, a_1 \rangle \mid c_2\langle \rangle.$$

The translation selects one thread (in this case the first one) to continue. In fact, the original guarded process $P$ is assigned to the first coordinator because it is encoded as $\llbracket P \rrbracket_{c_1,a_1,j_1}$. The messages $j_1$ and $j_2$ makes coordinators part of the same transaction, as described previously. The last message $c_2\langle \rangle$ indicates that the second thread is finished by notifying its coordinator to start the commit protocol with vote commit.

In the previous description of the encoding we just focused on a few forms of processes (for instance, we have discarded all processes containing join patterns with more than two messages). Nevertheless, it can be shown that this syntactical restriction, called the class of *canonical processes*, does not change the expressiveness of cJoin (formal details can be found in (Melgratti 2005)). For the sake of clarity, we report below the formal statement of the results ensuring the correctness and completeness of our encoding. We will use $\rightarrow_J$ and $\rightarrow_{cJ}$ to distinguish reductions in Join from those in cJoin. Moreover, the notion of process equivalence we use relies on barbs defined as follows:

**Definition 5.1 (Barb).** The observation predicate $\downarrow_x$, also known as the *strong barb*, detects whether a process emits on some free name $x$:

$$P \downarrow_x \quad \text{iff} \quad \exists P', \vec{u} : P \equiv \text{\bf def } D \text{ \bf in } P'|x\langle \vec{u} \rangle \text{ and } x \notin dn(D).$$

Note, the processes *abort* and [_ : _] have no strong barbs. Moreover, merge names are part of the defined names of a process, and hence not observable.

**Lemma 5.2.** For any canonical flat process $P$ s.t. $\vdash P : \square_1$ we have $\forall x.P \downarrow_x \Leftrightarrow [\![P]\!] \downarrow_x$.

Correspondence results assume the following property about the commit protocol: compensations corresponding to every coordinator in a transaction are released when at least one coordinator is required to abort while all continuations are released when every coordinator in the transaction is required to commit. Otherwise, none coordinator finishes, and none compensations nor continuations are released. It has been shown in (Melgratti 2005) that the D2PC used in our encoding satisfies above condition.

**Theorem 5.3 (Correctness, part 1).** Let $P$ be a canonical flat process and $\vdash P : \square_1$. If $P \to^*_{cJ} P'$ either $P' \equiv P$ or the following two conditions hold:

1. $P' \equiv \mathbf{def}\ D'\ \mathbf{in}\ M' \mid \Pi_{i\in1..n'}\mathcal{N}_i$, where $\mathcal{N}_i$ are cJoin transactions,
2. $\exists Q$ s.t. $[\![P]\!] \to^*_J Q$ and $Q \equiv (\mathbf{def}\ [\![D']\!]\ \mathbf{in}\ [\![M']\!]|\Pi_{i\in1..n'}R_i)|\mathbf{def}\ D_g\ \mathbf{in}\ 0$, where each $R_i$ is the *standard Join negotiation*[†] associated to $\mathcal{N}_i$ and $D_g$ collects garbage definitions corresponding to instances of the commit protocol that have terminated.

*Proof sketch.* The proof follows by case analysis on $P$. Note that $P$ cannot be of the form $[P' : Q']$ because $\not\vdash [P' : Q'] : \square_1$. If $P$ has no local definitions (i.e., $P \not\equiv \mathbf{def}\ D\ \mathbf{in}\ M$), then $P$ is either *abort*, the inert process 0, or the parallel composition of messages (containing only free names because there are no local definitions). In all three cases, $P' = P$. Last case is when $P$ contains local definitions, i.e., $P \equiv \mathbf{def}\ D\ \mathbf{in}\ M$. For this case, we show by induction on the length of the derivation of $P \to^*_{cJ} P'$ that the conditions 1 and 2 of the thesis hold (see appendix). $\square$

**Theorem 5.4 (Correctness, part 2).** Let $P$ be a canonical flat process and $\vdash P : \square_1$. If $P \to^*_{cJ} P'$ and $\vdash P' : \square_1$, then $\exists Q$ s.t. $[\![P]\!] \to^*_J Q$ and $\forall x.P' \downarrow_x \Rightarrow Q \downarrow_x$.

*Proof sketch.* By Theorem 5.3, we know that either $P' \equiv P$ or the following two conditions hold:

1.  $P' \equiv \mathbf{def}\ D'\ \mathbf{in}\ M' \mid \Pi_{i\in1..h'}\mathcal{N}_i$, where $\mathcal{N}_i$ are cJoin transactions,
2.  $\exists Q$ s.t. $[\![P]\!] \to^*_J Q$ and $Q \equiv (\mathbf{def}\ [\![D']\!]\ \mathbf{in}\ [\![M']\!]|\Pi_{i\in1..h'}R_i)|\mathbf{def}\ D_g\ \mathbf{in}\ 0$.

If $P' \equiv P$, then the thesis trivially follows by taking $Q = [\![P]\!]$ and applying Lemma 5.2.

Otherwise, $P \equiv \mathbf{def}\ D\ \mathbf{in}\ M$ for some $D$ and $M$, and by $\vdash P' : \square_1$, it must be the case that $P' \equiv \mathbf{def}\ D'\ \mathbf{in}\ M'$ for some $D'$ and $M'$.

By property above, $\exists Q$ s.t. $[\![P]\!] \to^*_J Q = \mathbf{def}\ [\![D']\!]\ \mathbf{in}\ [\![M']\!] \mid \mathbf{def}\ D_g\ \mathbf{in}\ 0$. It is easy to notice that $\forall x : P' \downarrow_x \Rightarrow Q \downarrow_x$ because the encoding ensures that $fn(M') = fn([\![M']\!])$ and $dn([\![D']\!]) \cap fn(M) = \varnothing$. $\square$

**Theorem 5.5 (Completeness).** Let $P$ be a canonical flat process and $\vdash P : \square_1$. If $[\![P]\!] \to^*_J Q$, then $P \to^*_{cJ} P'$ and $\forall x : norm(Q) \downarrow_x \Rightarrow P' \downarrow_x$, where $norm(Q)$ denotes the process obtained

---

[†] A standard Join negotiation is the Join counterpart of a cJoin transaction. It basically consists of the set of transaction coordinators belonging to the same transaction and of all the execution threads associated to those coordinators. Its formal definition can be found in (Melgratti 2005).

from $Q$ by finishing the execution of all instances of the commit protocol that can reach an agreement.

*Proof.* We proceed by case analysis on the structure of $P$. Since $\vdash P : \square_1$, then $P \equiv [P' : Q']$. When $P$ has no local definitions, then it is the parallel composition of messages on free ports, the inert process 0 and *abort*. For any of these cases, it holds that $[\![P]\!]$ does not have any definition, and therefore $[\![P]\!]$ cannot reduce. The only possibility is $Q = norm(Q) = [\![P]\!]$, which trivially satisfies $\forall x : Q \downarrow_x \Rightarrow [\![P]\!] \downarrow_x$. If $P \equiv \textbf{def } D \textbf{ in } M$, then we show that the following three conditions hold:

1. $Q \equiv \textbf{def } [\![D']\!] \textbf{ in } [\![M_1']\!] \mid \Pi_{i \in 1..u} R_i' \mid \Pi_{k \in 1..f} T_k' \mid \textbf{def } D_g \textbf{ in } 0$, where $R_i'$ are unfinished Join negotiations (i.e., some transactional thread has not finished), while $T_k'$ are finished negotiations, with $norm(\Pi_{k \in 1..f} T_k') \equiv [\![M_2]\!] \mid \textbf{def } D_c \textbf{ in } 0$.
2. $P \rightarrow_{cJ}^* P' \equiv \textbf{def } D' \textbf{ in } M_1' \mid M_2 \mid \Pi_{i \in 1..u} \mathcal{N}_i$, where $\mathcal{N}_i$ is a standard cJoin transaction corresponding to $R_i'$.
3. $norm(Q) \equiv \textbf{def } [\![D']\!] \textbf{ in } [\![M_1'|M_2]\!] \mid norm(\Pi_{i \in 1..u} R_i') \mid \textbf{def } D_g' \textbf{ in } 0$.

Above conditions are proved by induction on the length of the derivation $[\![P]\!] \rightarrow_J^* Q$ (see appendix).

Finally, condition $\forall x : norm(Q) \downarrow_x \Rightarrow P' \downarrow_x$ immediately follows from conditions (2) and (3).

$\square$

We chose not to report here the full details about the encoding since the formal definition gets quite complex because of the several alternatives to be considered in the most general case. These alternatives came from the fact that the encoding of a guarded process (in addition to what was explained before) may depend on the type of the received names, in particular whether they are transactional or not. However, the increased complexity of the notation makes it heavy with no evident conceptual benefit. The interested reader can found the full details spelled out in (Bruni *et al.* 2003; Melgratti 2005).

## 6. t-JoCaml

We take advantage of the transactional mechanism of cJoin to extend a programming language with transaction primitives. We have chosen JoCaml (Conchon and Le Fessant 1999), one of the available implementations of Join. We start this section by giving an overview of JoCaml. Then, we describe the transactional extension we propose, called *transactional* JoCaml (t-JoCaml), and finally, we sketch the main aspects of t-JoCaml implementation.

### 6.1. *JoCaml*

JoCaml adds Join primitives to *Objective Caml* (Ocaml), which is a functional language with support for object-oriented and imperative paradigms. JoCaml provides three main abstractions: *process*, *channels* and *join-patterns*. Processes represent communication and synchronisation tasks. Basic processes are asynchronous messages, while complex processes

are obtained by composing expressions and concurrent processes. Channels are JoCaml abstractions corresponding to Join names. There are two different kind of channels: *synchronous* and *asynchronous*. Channels are defined as follows:

$$\textbf{let def} \quad name[!](args) = P(args); ;$$

The above definition creates a channel (named *name*) and a receiver for it, which will execute the *guarded process* $P$ every time it receives a message. Any channel may be defined either as asynchronous, when its name is suffixed with the symbol !, or as synchronous, otherwise. Synchronous names must return a value, i.e., $P$ must explicitly define the return value $v$ by using the sentence **reply** $v$. Finally, join-patterns define several channels at the same time and state a synchronisation among them: the guarded process may be activated only when all channels have pending messages. The following is a possible join pattern definition:

$$\textbf{let def} \quad a!(x) \mid b!(y) = P(x, y)$$
$$\textbf{or} \quad a!(x) \mid c!(z) = Q(x, z)$$
$$; ;$$

The process above introduces three new asynchronous ports, namely $a$, $b$ and $c$. Like join processes, the guarded process $P$ (depending on variables $x$ and $y$) can be activated only when both $a$ and $b$ have pending messages. Similarly, $Q(x, z)$ can be activated when both $a$ and $c$ have pending messages. Moreover, when both rules are enabled, the selection is unspecified.

Processes can also create fresh ports dynamically. Consider the following program:

$$\textbf{let def} \; new\_process() =$$
$$\textbf{let def} \quad a!() \mid b!() = P$$
$$\textbf{or} \quad a!() \mid c!() = Q$$
$$\textbf{in reply} \; a, b, c$$
$$; ;$$

It declares a synchronous port *new_process* (i.e., an ordinary function) that, when called, creates a new process defining three fresh ports $a$, $b$ and $c$. The caller is given back the names of the created ports (by clause **in reply** $a, b, c$).

### 6.2. Transactions for JoCaml

In order to add transactions to JoCaml, we extend its syntax by allowing the definition of a compensable transaction, abort decision, and merge definitions.

6.2.1. *t-JoCaml syntax:* as already mentioned, we added two new forms of processes to JoCaml: transactional processes and abort. A transactional process is written as **let trans** $P$ **cmp** $Q$, where $P$ is an ordinary JoCaml program and $Q$ may be either an ordinary JoCaml process or a transactional one. For instance, the client component in

Figure 6, can be written as:

$$\textbf{let trans def} \quad offer\,!(rate, k) = k(card)$$
$$\textbf{or} \quad offer\,!(rate, k) = \textbf{abort}$$
$$\textbf{in} \; HotelReq(data, offer)$$
$$\textbf{cmp} \; Q$$
$$;;$$

Note the straightforward correspondence with the cJoin definition in Figure 6. Similarly to cJoin, transactional processes in t-JoCaml may decide to abort the execution of a transaction by using the new primitive **abort**.

Merge patterns are defined by writing the keyword **board** in front of the corresponding join patterns. Then, merge definition for the booking trip system in Figure 6, can be written as below:

$$\textbf{let board} \; HotelSrv\,!(r) \mid HotelReq\,!(d, k) = r(d, k) \; \textbf{in} \ldots$$

As the corresponding Join definition, the above sentence introduces two merge ports *HotelSrv* and *HotelReq*, with the guarded process $r(d, k)$, which is required to be an ordinary JoCaml process (i.e., without transaction primitives).

6.2.2. *Extending JoCaml compiler:* as far as the compiler implementation is concerned, we translate syntactically JoCaml programs with transactions into ordinary JoCaml code. We do this by reusing the parsing phase of the JoCaml distribution by slightly modifying the lexer and the parser in order to recognise processes built with the primitives **trans**, **comp**, **abort** and **board**. After the construction of the parse tree, we generate a new file containing the corresponding JoCaml source code that uses the encoding presented in Section 5. The implementation is available at `http://www.di.unipi.it/~melgratt/cjoin`.

The main limitation of the current prototype version is that it cannot handle the compilation of separate units. This restriction relies on the fact that the translation is parametric on the types of free names (i.e., whether they are ordinary or merge channels). Current translation of a program assumes that the typing environment is initially empty and it is updated when new port definitions are introduced by the program. Consequently, our prototype is unable to handle merge definitions introduced by different files. This constraint could be overcome by adding a new primitive for importing declarations of merge ports explicitly.

## 7. Big-step semantics and serializability

In this section, we introduce an alternative definition for the semantics of cJoin that allows us to reason about transactional computations at different levels of abstraction. The big-step semantics is intended to single out those computations of a system that are not transient or, in other words, to describe the evolution of a system through states that do not contain active transactions.

**Example 8.** Consider the cJoin process $P \equiv$ **def** D **in** $c\langle a\rangle | c\langle b\rangle$ with D defined as follows:

$$
\begin{aligned}
\mathsf{D} \equiv \quad & a\langle x\rangle | b\langle y\rangle \blacktriangleright ok\langle\rangle \\
& \wedge a\langle x\rangle | b\langle y\rangle \blacktriangleright abort \\
& \wedge c\langle x\rangle \triangleright [\mathbf{def}\ z \triangleright 0\ \mathbf{in}\ x\langle z\rangle : Q].
\end{aligned}
$$

The process $P$ may evolve as follows:

$$
\begin{aligned}
P \ \to\ & P_1 \equiv \mathbf{def}\ \mathsf{D}\ \mathbf{in}\ [\mathbf{def}\ z \triangleright 0\ \mathbf{in}\ a\langle z\rangle : Q\{^a/_x\}] \mid c\langle b\rangle \\
\to\ & P_2 \equiv \mathbf{def}\ \mathsf{D}\ \mathbf{in}\ [\mathbf{def}\ z \triangleright 0\ \mathbf{in}\ a\langle z\rangle : Q\{^a/_x\}] \mid [\mathbf{def}\ z \triangleright 0\ \mathbf{in}\ b\langle z\rangle : Q\{^b/_x\}] \\
\to\ & P_3 \equiv \mathbf{def}\ \mathsf{D}\ \mathbf{in}\ [\mathbf{def}\ z \triangleright 0 \wedge z' \triangleright 0\ \mathbf{in}\ ok\langle\rangle : Q\{^a/_x\} | Q\{^b/_x\}] \\
\to\ & P_4 \equiv \mathbf{def}\ \mathsf{D}\ \mathbf{in}\ ok\langle\rangle.
\end{aligned}
\tag{1}
$$

Analogously,

$$
\begin{aligned}
P \to^* P_2 \ \to\ & P_5 \equiv \mathbf{def}\ \mathsf{D}\ \mathbf{in}\ [\mathbf{def}\ z \triangleright 0 \wedge z' \triangleright 0\ \mathbf{in}\ abort : Q\{^a/_x\} | Q\{^b/_x\}] \\
\to\ & P_6 \equiv \mathbf{def}\ \mathsf{D}\ \mathbf{in}\ Q\{^a/_x\} | Q\{^b/_x\}.
\end{aligned}
$$

Assuming $Q$ does not have any active transaction, the computations above are the only two possible evolutions of $P$ to non-transient states, i.e., states that do not contain running transactions. We are aimed at defining a reduction relation $\twoheadrightarrow$ in which any multi-party transaction is described as a single computation step that fetches the messages needed to initiate all cooperating transactions and produces the processes released at commit or abort. In this example, we expect two big-step reductions for $P$, each of them describing one of the possible executions of the multi-party transaction, namely: $P \twoheadrightarrow P_4$ and $P \twoheadrightarrow P_6$.

We define the big-step reduction relation for a particular class of processes, called *shallow*. Shallow processes are given in terms of a syntactic restriction that imposes a particular discipline for activating transactions. We start by introducing the class of shallow processes.

**Definition 7.1 (Nesting level).** The *nesting level* (or just *nesting*) of $P$, written $nest(P)$, is defined by:

$$
\begin{array}{ll}
nest(0) = nest(abort) = nest(x\langle y\rangle) = 0 & nest([P : Q]) = nest(P) + 1 \\
nest(\mathbf{def}\ D\ \mathbf{in}\ P) = nest(P) & nest(P \mid Q) = \mathsf{max}\{nest(P), nest(Q)\}.
\end{array}
$$

We remark that $nest(P)$ counts only the nesting level of the active processes independently from the nested transactions that may appear in definitions or compensations. Consider $P \equiv [x\langle y\rangle : 0]$ and $Q \equiv [P : 0]$, then $nest(P) = 1$ and $nest(Q) = nest(P) + 1 = 2$. Contrastingly, $nest([0 : Q]) = 1$ because $[0 : Q]$ is a transaction that has no active sub-transactions (note that the compensation $Q$ is frozen, i.e., inactive).

**Definition 7.2 (Shallow and stable processes).** A basic definition $D$ is a *shallow definition* if it has one of the following forms:

1. $D = J \triangleright P$, where $nest(P) = 0$ or $P = [R : Q] \wedge nest(R|Q) = 0$.
2. $D = J \blacktriangleright P$ and $nest(P) = 0$.

A process $P$ is *shallow* if any basic definition in $P$ is shallow. Moreover, we call a shallow process $P$ *stable* iff $nest(P) = 0$.

With $P$ and $Q$ as defined above, the process $R \equiv \textbf{def}\ x\langle y \rangle \triangleright Q\ \textbf{in}\ x\langle a \rangle$ is stable (i.e., $nest(R) = 0$) because it does not have any active transaction (independently from the fact that it may start a transaction in the future).

Shallowness is a constraint for the syntax of basic definitions contained by a process. Condition 1 in Definition 7.2 ensures that the firing of a basic definition increases the height of the nesting structure by at most one level, i.e., a basic definition produces either a stable process or an activate transaction without any active sub-transaction. Condition 2 forbids the creation of sub-transactions while merging. We remark that shallowness does not impose any constraint on the nesting level of active transactions. For instance, let us consider $Q \equiv [[x\langle y \rangle : 0] : 0]$ and $R \equiv \textbf{def}\ x\langle y \rangle \triangleright Q\ \textbf{in}\ x\langle a \rangle$. The process $Q$ trivially satisfies shallowness condition because it does not have any definition. On the contrary, the process $R$ is not shallow because its unique definition $x\langle y \rangle \triangleright Q$ does not satisfy condition 1. In fact, $nest([x\langle y \rangle : 0]|0) = 1 \neq 0$.

We highlight that any flat process is also shallow. It can be easily seen that shallowness is preserved by reductions, while in general this is neither the case for the nesting of a process nor for stability (i.e., any shallow process always reduces to shallow processes, while some stable processes may reduce to non-stable processes).

Moreover, it can be shown that any non-shallow definition can be encoded into a shallow definition. Shallowness forbids definitions like $D_0 \equiv x\langle y \rangle \triangleright [[y\langle x \rangle : 0] : 0]$ and, more generally, $D_1 \equiv J_1 \triangleright P \mid [P_1 : Q]$ and $D_2 \equiv J_2 \triangleright [[P_1 : Q_1] : Q]$, they however can be encoded as shallow definitions by using new local ports. The cases above can be rewritten as $D'_0 \equiv x\langle y \rangle \triangleright [\textbf{def}\ z\langle \rangle \triangleright [y\langle x \rangle : 0]\ \textbf{in}\ z\langle \rangle : 0]$, $D'_1 \equiv J_1 \triangleright z\langle \rangle|P\ \wedge\ z\langle \rangle \triangleright [P_1 : Q]$ and $D'_2 \equiv J_2 \triangleright [\textbf{def}\ z\langle \rangle \triangleright [P_1 : Q_1]\ \textbf{in}\ z\langle \rangle : Q]$ with $z$ fresh. Note that the three new definitions are shallow when $P$, $P_1$, $Q$, $Q_1$ are stable. Then, $\textbf{def}\ D'_1\ \textbf{in}\ x\langle a \rangle$ is a shallow process and reduces in two steps to $\textbf{def}\ D'_1\ \textbf{in}\ [\textbf{def}\ z\langle \rangle \triangleright [a\langle x \rangle : 0]\ \textbf{in}\ [a\langle x \rangle : 0] : 0]$, which is shallow and contains nested transactions.

In the following, $\mathcal{P}$ and $\mathcal{Q}$ will denote shallow processes, $\mathcal{D}$ a shallow definition, $\mathcal{S}$ a stable process, and $\mathcal{B}$ a shallow definition containing just merge rules. We abbreviate $\textbf{def}\ \mathcal{D}\ \textbf{in}\ \mathcal{P}$ as $\mathcal{D} \vdash \mathcal{P}$, and $\vdash \mathcal{P}$ as $\mathcal{P}$. Terms are considered up-to structural equivalence generated by closure w.r.t. the equations for the associativity and commutativity of $\mid$ and $\wedge$, 0 the unit for $\mid$, and

$$\mathcal{D} \vdash (\mathcal{P} \mid \textbf{def}\ \mathcal{D}'\ \textbf{in}\ \mathcal{Q}) = \mathcal{D} \wedge \mathcal{D}'\sigma_{dn} \vdash \mathcal{P} \mid \mathcal{Q}\sigma_{dn}$$
$$range(\sigma_{dn}) \cap (fn(\mathcal{D})\ \cup\ fn(\mathcal{P})\ \cup\ fn(\textbf{def}\ \mathcal{D}'\ \textbf{in}\ \mathcal{Q})) = \varnothing.$$

The big-step reduction relation $\twoheadrightarrow$ is given by the inference rules in Figure 12, over stable processes. The relation $\twoheadrightarrow$ is defined in terms of the auxiliary relation $\rightarrow$ over shallow processes. Rule SERIALIZABLE singles out as big steps those computation from stable states to stable states. Note that computation steps can be composed in parallel (PAR) and sequentially (SEQ), even with idle transitions (IDLE). Rule GLOBAL FIRING, abbreviated GF, corresponds to the firing of an ordinary definition in a top-level process. Instead LOCAL FIRING states possible internal transitions of a running transaction. LOCAL FIRING represents suitable sub-transactions as ordinary transitions at an abstract level. In fact, the computations occurring at a lower level in the nesting hierarchy (premise

(PAR)
$$\frac{\mathcal{D} \vdash \mathcal{P} \ \twoheadrightarrow\ \mathcal{D} \vdash \mathcal{P}' \quad \mathcal{D} \vdash \mathcal{Q} \ \twoheadrightarrow\ \mathcal{D} \vdash \mathcal{Q}'}{\mathcal{D} \vdash \mathcal{P} \mid \mathcal{Q} \ \twoheadrightarrow\ \mathcal{D} \vdash \mathcal{P}' \mid \mathcal{Q}'}$$

(SEQ)
$$\frac{\mathcal{D} \vdash \mathcal{P} \ \twoheadrightarrow\ \mathcal{D} \vdash \mathcal{P}'' \quad \mathcal{D} \vdash \mathcal{P}'' \ \twoheadrightarrow\ \mathcal{D} \vdash \mathcal{P}'}{\mathcal{D} \vdash \mathcal{P} \ \twoheadrightarrow\ \mathcal{D} \vdash \mathcal{P}'}$$

(GLOBAL FIRING)
$$\mathcal{D} \wedge J \triangleright \mathcal{P} \vdash J\sigma \ \twoheadrightarrow\ \mathcal{D} \wedge J \triangleright \mathcal{P} \vdash \mathcal{P}\sigma$$

(LOCAL FIRING)
$$\frac{\widetilde{\mathcal{B}} \vdash \mathcal{S} \ \twoheadrightarrow\ \widetilde{\mathcal{B}} \vdash \mathcal{S}'}{\mathcal{D} \wedge \mathcal{B} \vdash [\mathcal{S} : \mathcal{Q}] \ \twoheadrightarrow\ \mathcal{D} \wedge \mathcal{B} \vdash [\mathcal{S}' : \mathcal{Q}]}$$

(MERGE)
$$\mathcal{D} \wedge \Pi_i J_i \blacktriangleright \mathcal{S} \ \vdash \ \Pi_i[\mathcal{D}_i \vdash J_i\sigma \mid \mathcal{S}_i : \mathcal{Q}_i] \ \twoheadrightarrow\ \mathcal{D} \wedge \Pi_i J_i \blacktriangleright \mathcal{S} \ \vdash \ [\bigwedge_i \mathcal{D}_i \vdash (\Pi_i \mathcal{S}_i) \mid \mathcal{S}\sigma : \Pi_i \mathcal{Q}_i]$$

(LOCAL COMMIT)         (ABORT)          (IDLE)
$$\mathcal{D} \vdash [M \mid \mathcal{D}' \vdash 0 : \mathcal{S}] \ \twoheadrightarrow\ \mathcal{D} \vdash M \qquad \mathcal{D} \vdash [abort \mid \mathcal{P} : \mathcal{S}] \ \twoheadrightarrow\ \mathcal{D} \vdash \mathcal{S} \qquad \mathcal{D} \vdash \mathcal{P} \ \twoheadrightarrow\ \mathcal{D} \vdash \mathcal{P}$$

(SERIALIZABLE)
$$\frac{\mathcal{D} \vdash \mathcal{S} \ \twoheadrightarrow\ \mathcal{D} \vdash \mathcal{S}'}{\mathcal{D} \vdash \mathcal{S} \ \twoheadrightarrow\!\!\!\twoheadrightarrow\ \mathcal{D} \vdash \mathcal{S}'}$$

Fig. 12. Big-step semantics of cJoin.

of LOCAL FIRING) that are relevant to its containing transaction are those relating stable processes, i.e., $\mathcal{S}$ and $\mathcal{S}'$. A transaction has available, in addition to its own definitions, the merge definitions introduced by its parent. In fact, a merge definition applied to a single transaction behaves as an ordinary rule but defined in a global scope. The operator $\widetilde{\underline{\ }}$ transforms merge definitions in ordinary ones: $\widetilde{J \blacktriangleright \mathcal{P}} = J \triangleright \mathcal{P}$ and $\widetilde{\mathcal{B} \wedge \mathcal{B}'} = \widetilde{\mathcal{B}} \wedge \widetilde{\mathcal{B}'}$.

Rules LOCAL COMMIT (abbreviated LC) and ABORT handle the termination of a transaction, whereas MERGE describes the interaction among sibling transactions. This time, transactions can be joined only if they do not contain running transactions.

**Example 9.** Consider the process $P$ introduced in Example 8. Processes $P$, $P_4$ and $P_6$ are stable and it can be easily checked that $P \twoheadrightarrow\!\!\!\twoheadrightarrow P_4$ and $P \twoheadrightarrow\!\!\!\twoheadrightarrow P_6$ as expected. A proof for the big-step reduction corresponding to the small-step reduction shown in Equation (1) of Example 8 can be built as follows: first, the small step $P \rightarrow P_1$ corresponds to the following proof:

$$\frac{\mathsf{D} \vdash c\langle a \rangle \ \twoheadrightarrow\ \mathsf{D} \vdash [\textbf{def } z \triangleright 0 \ \textbf{in } a\langle z \rangle : Q\{^a/_x\}] \ \text{GF} \quad \mathsf{D} \vdash c\langle b \rangle \ \twoheadrightarrow\ \mathsf{D} \vdash c\langle b \rangle \ \text{IDLE}}{P = \mathsf{D} \vdash c\langle a \rangle \mid c\langle b \rangle \ \twoheadrightarrow\ P_1 = \mathsf{D} \vdash [\textbf{def } z \triangleright 0 \ \textbf{in } a\langle z \rangle : Q\{^a/_x\}] \mid c\langle b \rangle} \ \text{PAR.}$$

The proof for $P_1 \twoheadrightarrow P_2$ is built analogously. Then, the proof for $P \twoheadrightarrow\!\!\!\twoheadrightarrow P_4$ can be completed as follows:

$$\frac{\dfrac{\dfrac{\vdots \quad \vdots}{\dfrac{P \twoheadrightarrow P_1 \quad P_1 \twoheadrightarrow P_2}{P \twoheadrightarrow P_2} \ \text{SEQ} \quad P_2 \twoheadrightarrow P_3 \ \ \text{MERGE}}}{\dfrac{P \twoheadrightarrow P_3}{\dfrac{P \twoheadrightarrow P_4}{P \twoheadrightarrow\!\!\!\twoheadrightarrow P_4} \ \text{SERIALIZABLE}} \ \text{SEQ}} \ \text{SEQ} \quad P_3 \twoheadrightarrow P_4 \ \ \text{GLOBAL COMMIT}}$$

Fig. 13. Big-step reductions for the Example 8.

Consider now the process $P|P$. Figure 13 shows the possible evolutions of $P|P$. The graph illustrates how the result of a computation involving the (possible interleaved) execution of several multi-party transactions, e.g., $P|P \twoheadrightarrow P_4|P_6$, can be also obtained by executing one transaction at a time. For $P|P \twoheadrightarrow P_4|P_6$ we can sequentially compute either $P|P \twoheadrightarrow P|P_4 \twoheadrightarrow P_4|P_6$ or $P|P \twoheadrightarrow P|P_6 \twoheadrightarrow P_4|P_6$.

It is in this sense that the big-step reduction relation enforces serializability and allows us to analyse the behaviour of a set of interacting transactions independently from the rest of the system. Moreover, when considering nested transactions, the transactions completed at a particular level of nesting can be treated as ordinary transitions at the upper level.

The remaining of this section is devoted to show the correspondence between both semantics for shallow processes, which ensures that small-step reductions are serializable for shallow processes.

The next auxiliary result shows that any derivation in cJoin starting from a shallow process without nested transactions has an equivalent cJoin reduction that only merges transactions without ongoing sub-transactions.

**Proposition 7.3.** If $\mathcal{D}, \mathcal{S}, \otimes_k [\mathcal{S}'_k : \mathcal{S}''_k] \to^* \mathcal{D}', \otimes_i [\mathcal{P}_i : \mathcal{S}_i] \to \mathcal{D}', [\mathcal{P} : \|_i \mathcal{S}_i]$, then there exists a derivation $\mathcal{D}, \mathcal{S}, \otimes_k [\mathcal{S}'_k : \mathcal{S}''_k] \to^* \mathcal{D}', \otimes_i [\mathcal{S}'''_i : \mathcal{S}_i] \to \mathcal{D}', [\mathcal{S}' : \|_i \mathcal{S}_i] \to^* \mathcal{D}', [\mathcal{P} : \|_i \mathcal{S}_i]$.

*Proof.* By induction on the length of the derivation $\mathcal{D}, \mathcal{S}, \otimes_k [\mathcal{S}'_k : \mathcal{S}''_k] \to^n \mathcal{D}', \otimes_i [\mathcal{P}_i : \mathcal{S}_i]$. The base case follows immediately by taking $\mathcal{S} = \varnothing$ and $k = i$. The inductive step follows by noting first that any reduction $\mathcal{D}, \mathcal{Q} \to \mathcal{D}', x\langle \vec{v} \rangle, \mathcal{Q}'$ implies either (i) $\mathcal{Q} = x\langle \vec{v} \rangle | \mathcal{Q}''$ or (ii) $\mathcal{Q} = [\mathcal{S}_1 : \mathcal{S}_2] | \mathcal{Q}''$ and $[\mathcal{S}_1 : \mathcal{S}_2] \to \mathcal{D}'', x\langle \vec{v} \rangle | \mathcal{S}_0$ (this can be shown by case analysis on the applied rule). This property tell us that the generation of a message does not imply the generation of a new transaction. Then, note that the only possibility for the last reduction $\mathcal{D}', \otimes_i [\mathcal{P}_i : \mathcal{S}_i] \to \mathcal{D}', [\mathcal{P} : \|_i \mathcal{S}_i]$ is due to the application of a merge reaction involving all transactions. Therefore, any $[\mathcal{P}_i : \mathcal{S}_i]$ is such that $\mathcal{P}_i = M_i | \mathcal{P}'_i$. If all $\mathcal{P}_i$ are stable we are done. Otherwise, we proceed by noting that any $[\mathcal{P}_i : \mathcal{S}_i]$ has been generated from the elements of the original solution $\mathcal{D}, \mathcal{S}, \otimes_k [\mathcal{S}'_k : \mathcal{S}''_k]$. The interesting case is when the whole reduction cannot be divided into sequences like $\mathcal{D}, \mathcal{S}, \otimes_k [\mathcal{S}'_k : \mathcal{S}''_k] \to^* \mathcal{D}'', \mathcal{S}'', \otimes_h [\mathcal{P}_h : \mathcal{P}_h] \to^* \mathcal{D}', \otimes_i [\mathcal{P}_i : \mathcal{S}_i]$ (these cases can be handled by using inductive hypothesis). Therefore, we note that there exists a partition of $\mathcal{S}, \otimes_k [\mathcal{S}'_k : \mathcal{S}''_k]$ s.t. for each $i$ there exists some $\mathcal{I}$ in the partition and $\mathcal{D}, \mathcal{I} \to^* \mathcal{D}'', [\mathcal{P}_i : \mathcal{S}_i]$. By using inductive hypothesis we can build $\mathcal{D}, \mathcal{I} \to^* \mathcal{D}'', \otimes_{k_i} [\mathcal{R}_{k_i} : \mathcal{R}'_{k_i}] \to \mathcal{D}'', [\mathcal{R} : \mathcal{R}'] \to^* \mathcal{D}'', [\mathcal{P}_i : \mathcal{S}_i]$. Since $\mathcal{P}_i = M_i | \mathcal{P}'_i$, we can show (by repeatedly using the property that ensures that transactions

$$\frac{\dfrac{\vdots \quad \text{GLOBAL FIRING,IDLE,PAR}}{\widetilde{\mathcal{B}} \vdash \mathcal{S}_k \;\rightarrowtail\; \widetilde{\mathcal{B}} \wedge \mathcal{D}'' \vdash [\mathcal{S}'' : \mathcal{P}'] \mid \mathcal{S}'_k} \qquad \dfrac{\widetilde{\mathcal{B}} \wedge \mathcal{D}'' \vdash [\mathcal{S}'' : \mathcal{P}'] \mid \mathcal{S}'_k \;\rightarrowtail\; \widetilde{\mathcal{B}} \vdash \mathcal{S}''_k \;\; \text{IND. HYP.}}{} }{\dfrac{\widetilde{\mathcal{B}} \vdash \mathcal{S}_k \;\rightarrowtail\; \widetilde{\mathcal{B}} \vdash \mathcal{S}''_k}{\mathcal{D}' \wedge \mathcal{B} \vdash [\mathcal{S}_k : \mathcal{P}_k] \;\rightarrowtail\; \mathcal{D}' \wedge \mathcal{B} \vdash [\mathcal{S}''_k : \mathcal{P}_k]} \;\text{LOCAL FIRING}} \;\text{SEQ}$$

Fig. 14. Proof sketch for Lemma 7.4.

are not generated together with messages) that $\mathcal{R}_{k_i} = M_i \mid \mathcal{R}_{k_i}$. Consequently, we can build a reduction that merges first all $[\mathcal{R}_{k_i} : \mathcal{R}'_{k_i}]$, and then reduces to the final configuration. □

In particular, serializable transactions can postpone the activation of each sub-transaction until all other cooperating sub-transactions needed to commit can be activated.

Next results state the correspondences between both semantics.

**Lemma 7.4.** $\mathcal{D}, \otimes_i[\mathcal{S}_i : \mathcal{P}_i], \mathcal{S} \rightarrow^* \mathcal{D}, \otimes_j[\mathcal{S}_j : \mathcal{P}'_j], \mathcal{S}'$ if and only if $\mathcal{D} \vdash \|_i[\mathcal{S}_i : \mathcal{P}_i] \mid \mathcal{S} \;\rightarrowtail\; \mathcal{D} \vdash \|_j[\mathcal{S}'_j : \mathcal{P}'_j] \mid \mathcal{S}'$.

*Proof.* ($\Rightarrow$) By induction on the length of the derivation. The base case corresponds to the IDLE axiom. The inductive step follows by case analysis of the first applied reduction. The interesting case is the application of RED, which presents two cases: (i) when it is applied at top level, i.e., producing a new transaction or a new stable process, and (ii) when it generates a sub-transaction in one of the existing transactions. First case is immediate by inductive hypothesis and SEQ. In the second case, there is a transaction $k$, i.e., $[\mathcal{S}_k : \mathcal{P}_k]$, that reduces to $[[\mathcal{S}' : \mathcal{P}'], \mathcal{S}'_k : \mathcal{P}_k]$. At this point, there are two possibilities. First, consider that $[[\mathcal{S}' : \mathcal{P}'], \mathcal{S}'_k : \mathcal{P}_k]$ reduces to $[\mathcal{S}''_k : \mathcal{P}_k]$, then it is possible to build the proof shown in Figure 14, and then the proof follows by using IDLE for the non-modified processes and inductive hypothesis and SEQ. The remaining possibility is when sub-transaction $[\mathcal{S}'' : \mathcal{P}']$ can finish only after the parent transaction is merged. Then by Proposition 7.3 an equivalent derivation $\mathcal{D}, \otimes_i[\mathcal{S}_i : \mathcal{P}_i], S \rightarrow^* \mathcal{D}, [\mathcal{S}'' : \|_i \mathcal{P}_i], S''' \rightarrow^* \mathcal{D}, [\mathcal{P} : \|_i \mathcal{P}_i], S''' \rightarrow^* \mathcal{D}, \otimes_j[\mathcal{S}_j : \mathcal{P}'_j], \mathcal{S}$ that merges only transactions not containing sub-transactions can be found. Then, proof follows by applying inductive hypothesis for both parts $\mathcal{D}, \otimes_i[\mathcal{S}_i : \mathcal{P}_i], S \rightarrow^* \mathcal{D}, [\mathcal{S}'' : \|_i \mathcal{P}_i], S'''$ and $\mathcal{D}, [\mathcal{S}'' : \|_i \mathcal{P}_i], S''' \rightarrow^* \mathcal{D}, [\mathcal{P} : \|_i \mathcal{P}_i], S''' \rightarrow^* \mathcal{D}, \otimes_j[\mathcal{S}_j : \mathcal{P}'_j], \mathcal{S}$.

($\Leftarrow$) By induction on the structure of the proof. □

**Theorem 7.5.** Let $\mathcal{S}, \mathcal{S}'$ be stable processes. Then $\mathcal{S} \rightarrow^* \mathcal{S}'$ iff $\vdash \mathcal{S} \;\rightarrowtail\; \vdash \mathcal{S}'$.

*Proof.* Immediate by Lemma 7.4. □

An informal explanation of the serializability result can be given by colouring transaction scopes and reductions as explained below. Let $\mathcal{S}, \mathcal{S}'$ be stable processes such that there exists $P_0, P_1, \ldots, P_k$ with $\mathcal{S} \equiv P_0 \rightarrow P_1 \rightarrow \ldots \rightarrow P_k \equiv \mathcal{S}'$. We traverse the computation backward, one reduction $P_i \rightarrow P_{i+1}$ at the time. If the reduction is originated from a commit or abort, then we assign it a fresh colour and use the same colour to paint the brackets of the corresponding transaction scope in $P_i$. If the reduction is originated from the merge of several transactions, then we paint the reduction and all the involved

transaction scopes in $P_i$ with the same colour as the merged transaction scope in $P_{i+1}$. If the reduction is an ordinary one, then we paint it with the same colour as the one of the immediately enclosing brackets of its target term in $P_{i+1}$ and we use the same colour to paint the immediately enclosing brackets of its source term in $P_i$, if any. At each step, we paint with the same colours as in $P_{i+1}$ all the transaction scopes that are not directly involved in the step. Then each colour can be viewed as representing activities of a distinct transaction, and we say that $c_1$ is a sub-colour of $c_2$ if the transaction associated with $c_1$ is a sub-transaction of the one associated with $c_2$. The serializability result essentially guarantees that another sequence of reductions $\mathcal{S} \equiv Q_0 \to Q_1 \to \dots \to Q_k \equiv \mathcal{S}'$ can be found, such that all steps of different colours are either contiguous or separated by steps of some sub-colour.

## 8. Related work and concluding remarks

We have proposed cJoin as a formal framework for designing and programming multi-party LRTs. Our calculus features name mobility, asynchronous communication and has a prototype implementation called t-JoCaml. We have included several examples that witness the flexibility of the calculus, together with a serializability result that hold for a wide range of processes, called shallow. The encodability of full cJoin in Join is an open issue, for which we have found no solution yet, because it would require the implementation of a scoping discipline for restricting communication over distributed processes. Nevertheless, in Sections 3 and 4 we have shown that the implemented fragment is valuable enough to model a large variety of frequently used patterns. Section 8.1 also witnesses that our proposal is quite original w.r.t. other ones found in the literature.

Transactions have been largely studied by the database community as the main mechanism for ensuring data consistency when executing concurrent sequences of operations (Eswaran *et al.* 1976). Many different models were proposed to meet the so-called ACID properties (i.e., all or nothing transactions) such as the *flat model* (Eswaran *et al.* 1976), *flat transactions with save-points* and *chained transactions* (Gray and Reuter 1993), *nested transactions* (Moss 1981) and the *multi-level model* (Lomet 1992; Schek and Weikum 1992; Weikum 1991) among others. These models are based on locking mechanisms that prevent concurrent transactions from accessing shared objects simultaneously (Bernstein *et al.* 1987; Fekete *et al.* 1994; Gray and Reuter 1993; Kohler 1981). Hence, the execution of a transaction may suffer considerable delays while waiting for others transactions to commit. Therefore, ACID transactions are regarded suitable only for handling transactions with short duration. Alternative models for long running transactions leave out ACID properties by relying on weaker notions of atomicity. The seminal proposal in this direction is Sagas (Garcia-Molina and Salem 1987), which introduces the model of multi-step transactions with programmed compensations described in Section 4.1. This model has been generalised by *Open Nested transactions* (Schek and Weikum 1992), in which compensable steps can be organised hierarchically. Several other models have appeared in the literature for allowing a flexible description of steps dependencies, such in (Elmagarmid *et al.* 1990; Hutchinson *et al.* 1988; Kaiser and Pu 1992).

As previously mentioned, the main goal of the above research line is to ensure consistent database updates. Differently, transactions in composition languages are aimed at coordinating atomic executions of independent activities. In this sense, composition languages are directly related to workflow management systems (wms) ( see (Georgakopoulos *et al.* 1995; Rusinkiewicz and Sheth 1995) for a detailed description of wms). Transactional execution of workflows has been an active research topic for wms community, as testified by several existing transactional wms, like Contracts (Reuter and Wächter 1992), meteor (Krishnakumar and Sheth 1995) or meteor2 (Kochut *et al.* 1996)). Recently, the works in (Bruni *et al.* 2005; Butler *et al.* 2002; Butler *et al.* 2005a; Butler *et al.* 2005b), further improved in (Bruni *et al.* 2012; Ripon 2008), have formally studied the semantics of orchestration languages with transactions. Differently from workflow community, this new research line is concerned about workflow systems as programming languages, and thus the focus is on giving precise, formal syntax, semantics and reasoning techniques for transaction primitives in orchestration.

Like previous approaches, cJoin is aimed at using programming language approach to study transactional composition. Nevertheless, it is targeted to the study of transactions for choreographies by adding transactions to a calculus for communicating processes. We devote the remaining part of this section to compare cJoin with other proposals from the literature that have similar aims.

### 8.1. *Language comparison*

In order to systematically analyse the proposals appeared in the literature, we first identify a set of aspects or choices for extending communicating processes calculi with transactions. Then, we present a comparison of the different approaches in terms of those selected features.

8.1.1. *Undoability.* A transactional mechanism provides a way for repairing the effects of the partial executions of aborted transactions. There are three main design options:

1. *Automatic rollback*: if a transaction $[P]$ aborts, then the scope $[\_]$ ensures that all the effects of the partial execution of $P$ are automatically removed from the state.
2. *Static programmable rollback*: a transactional process $P$ is associated statically with another process $Q$ such that $Q$ is activated whenever $P$ aborts. Programmers are responsible for writing $Q$ in such a way that the effects of the partial executions of $P$ are compensated for.
3. *Dynamic programmable rollback*: differently from static programmable rollback, compensations are built during execution. Programmers are responsible for describing how compensations change when transactions execute.
4. *Pre-committed compensation*: a LRT is divided into several steps. The successful execution of a step (i.e., it commits) may install ad hoc programs to be run only when the whole LRT aborts.

8.1.2. *Permeability.* Permeability refers to the degree of isolation provided by transactional scopes, i.e., whether messages can cross transactional scopes or not.

1. *Impermeable*: messages cannot flow across transactional boundaries, e.g., in $P \equiv x!z.0 \mid y?v.0 \mid [\ x?w.P \mid y!z.Q\ ]$ no communication is possible on $x$ and $y$.
2. *Permeable in input*: messages generated outside transactions can be received by transactional processes, e.g., process $P$ above can reduce by communicating over $x$ but not over $y$.
3. *Permeable in output*: transactional processes may send messages to receivers outside the transaction. In the previous example, $P$ can communicate on $y$ but not on $x$.
4. *Permeable*: messages freely flow across transactional boundaries. The previous process $P$ may communicate on both $x$ and $y$.
5. *Selective permeability*: messages may flow across transactional boundaries only when sent over channels that are in some particular class.

8.1.3. *Dynamicity.* Dynamicity characterizes the way in which the execution of a transactional scope may relate with the execution of other scopes.

1. *Static*: a static scope has no relation with other scopes, i.e., after being created it can neither affect nor be affected by the behaviour of other scopes.
2. *Joinable*: a scope is joinable if its abortion or commitment may condition or may be conditioned by the abortion or commitment of other scopes.
3. *Splittable*: a scope $[P]$ is splittable if it is possible to take a part of $P$ and run it as an independent scope.
4. *Dynamic*: a dynamic scope is both joinable and splittable.

8.1.4. *Naming.* The naming policy indicates the way in which transactional scopes are identified. In particular, they can be:

1. *Anonymous*: scopes have no explicit identification. This mean that transactional processes do not refer scopes explicitly.
2. *Named*: scopes have a name and processes refer to them explicitly. Usually a transaction is aborted by sending a message to its scope name. In addition, scope names may be *unique*, i.e., a name unequivocally identify a transactional scope; or *multiple*, i.e., a name may refer to several transactional scopes.

8.1.5. *Interaction model.* Transactional process calculi differ on the underlying interaction model. We use here a coarse-grain distinction in two main categories:

1. *Shared dataspaces:* processes communicate by writing to and reading from a shared blackboard.
2. *Message passing:* processes communicate by sending and receiving messages on specific ports or channels.

8.1.6. *Nesting.* Nesting relates to the capability of decomposing the execution of a transaction into a hierarchy of *sub-transactions*. In this scheme, any sub-transaction executes atomically and concurrently with respect to its parent and siblings, deciding freely to commit or abort. Nevertheless, if the parent aborts all its sub-transactions are also undone.

| | Undoability | Permeability | Dynamicity | Naming | Interaction | Nesting | Preemption |
|---|---|---|---|---|---|---|---|
| PLinda | *Automatic; Static* | *Input* | *Static* | *Anonym.* | *DataSpaces* | *No* | *May* |
| TSpaces | — | *Input* | *Static* | *Unique* | *DataSpaces* | *No* | — |
| TraLinda | *Automatic* | *Impermeable* | *Joinable* | *Anonym.* | *DataSpaces* | *No* | *May* |
| πt | *Static; Precomm.* | *Permeable* | *Static* | *Anonym.* | *Msg Passing* | *Yes* | *May* |
| webπ$_\infty$ | *Static* | *Permeable* | *Splittable* | *Multiple* | *Msg Passing* | *No* | *May* |
| RCCS | *Automatic* | *Permeable* | *Joinable* | *Anonym.* | *Msg Passing* | *No* | — |
| ρπ | *Automatic* | *Permeable* | *Joinable* | *Anonym.* | *Msg Passing* | *No* | — |
| dcπ | *Dynamic* | *Permeable* | *Static* | *Unique* | *Msg Passing* | *Yes* | *Must/ May* |
| *TransCCS* | *Automatic Static* | *Permeable* | *Joinable* | *Unique* | *Msg Passing* | *Yes* | — |
| cJoin | *Static* | *Selective* | *Joinable* | *Anonym.* | *Msg Passing* | *Yes* | *May* |

Fig. 15. Comparison of transactional process calculi.

8.1.7. *Preemption.* Following the classification in (Berry 1993), the abortion of an execution may take two different styles of preemption. Abortion may be (i) *may*-preemptive, i.e., a transaction may take an arbitrary number of internal computation steps before handling the abort condition, or (ii) *must*-preemptive, i.e., no further internal computation steps are allowed when a transaction reaches the abort (abortion is honoured immediately).

8.1.8. *Comparison.* We will use the previous seven categories to compare cJoin against the closest process calculi appeared in the literature: PLinda (Anderson and Shasha 1992), TSpaces (Busi and Zavattaro 2002), TraLinda (Bruni and Montanari 2004), πt (Bocchi *et al.* 2003), Webπ$_\infty$ (Lucchi and Mazzara 2004), RCCS (Danos and Krivine 2004), ρπ (Lanese *et al.* 2010a), dcπ (Vaz *et al.* 2008) and TransCCS (de Vries *et al.* 2010). Figure 15 summarises the features of all selected approaches. We remark that all considerations made for Webπ$_\infty$ are still valid for its timed version (Laneve and Zavattaro 2005).

Languages PLinda and TSpaces are aimed at providing a model for traditional serializable (i.e., ACID) transactions, hence they provide *input* permeability by allowing a

transaction to read or to consume data from a shared dataspace (i.e., the communication models represents a shared database). Differently, when the transaction produces a new datum, it is locked until the transaction commits and hence there is no *output* permeability. In both cases, transaction scopes are defined statically and cannot change dynamically. Transactions have no name in PLinda, while they are named in TSpaces. Names in TSpaces are included as a handy way for defining the semantics of the language, nevertheless, programmers do not need to be aware of them. Although, the syntax of TSpaces allows transaction names to be duplicated, its operational semantics ensures that each transactional name is treated as unique. None of those languages allows nesting (the syntax of TSpaces allows it but its semantics does not). Transactions in PLinda have an automatic, perfect rollback. Moreover, programmers have the possibility of specifying ad hoc programs to be run after rollback when a transaction aborts. Abortion is not considered in TSpaces.

TraLinda adds joinable transactions to a shared dataspace coordination language. Multi-way transactions of cJoin are analogous to TraLinda. Nevertheless, cJoin consider programmable compensations instead of perfect rollback, nested transactions instead of flat ones, and message passing communication instead of shared dataspace.

Both $\pi$t and Web$\pi_\infty$ are similar in spirit to cJoin. The main differences rely on the policies adopted for transactional scopes. While $\pi$t and Web$\pi_\infty$ allows transactional processes to freely interact with other processes, cJoin imposes a more strict policy: transactional processes may interact over selected channels only with transactional process but, in this case, they should reach the same decision (i.e., commit or abort). These differences come from the fact that scopes in $\pi$t and Web$\pi_\infty$ are permeable, while they are selective permeable and joinable in cJoin. All three languages provide a mechanism for programmable compensation (it is called fault handler in $\pi$t). In addition, $\pi$t provides a mechanism for undoing precommitted sub-transactions. Both $\pi$t and cJoin have nesting, while Web$\pi_\infty$ has not. Word nesting is used in Web$\pi_\infty$ to refer to splittable scopes.

RCCS and $\rho\pi$ provide transactions relying on a built-in distributed backtracking mechanism, which can achieve perfect rollback. Transactions are joinable in the sense that processes that have communicated are required to backtrack together. RCCS is an extension of CCS (Milner 1980), therefore, the underlying communication is just process synchronisation. On the contrary, $\rho\pi$ supports higher order communication (in fact, $\rho\pi$ extends RCCS to higher-order $\pi$). Differently from cJoin, RCCS and $\rho\pi$ do not support nesting, scopes are permeable, and transactions are automatically rolled-back. Abortion can be fired spontaneously (the execution of a transaction can be aborted at any time).

Compensations in cJoin are statically defined while they are dynamically built in dc$\pi$. The main idea behind dc$\pi$ is that any input prefix is associated with a compensation. Then, any time a process executes an input action, it also installs a compensation that will be activated if the corresponding transaction aborts. As shown in (Lanese *et al.* 2010b), dynamic compensations are more expressive than static compensations. Transactions in dc$\pi$ are completely permeable and static, like in $\pi$t and Web$\pi_\infty$. Hence, the commitment or abortion of one transaction does not affect the behaviour of the others. As for Web$\pi_\infty$, transactional scopes are named and their names are used for signalling abortion. Consequently, transaction in dc$\pi$ can be aborted internally or externally, while in cJoin

the abort condition can be reached only internally. Moreover, dcπ adopts may-preemption for handling abortion generated externally and must-preemption for internal abort. On the contrary, aborts are only internal and may-preemptive in cJoin. The main reason for this choice is that a transaction in cJoin can be the consequence of merging several independent, possible distributed transactions. Thus, the implementation of a must-preemptive abortion would be problematic in a distributed setting without central coordination.

The only joinable mechanism for transactions that we are aware of is the one proposed in TransCCS. As in cJoin, a transaction in TransCCS can be merged with other transactions during execution. Merging is completely symmetric in cJoin, i.e., the abort of a merged transaction releases the compensations corresponding to all original transactions. Differently, the merging of transactions in TransCCS generates a nested transaction (i.e., one transaction is included as sub-transaction of the other). Then, the abortion of one transaction in this hierarchy releases the compensation of the original transaction and automatically rolls back the original state of the transactions that has been included as a sub-transaction. Finally, abort is spontaneous in TransCCS.

We remark that other interesting approaches such as Pike (Chothia and Duggan 2004) and Transactional Linda (Jagannathan and Vitek 2004) have been left out from our discussion, since they are parametric frameworks in which the behaviour of a transaction does not rely on language primitives. Roughly, transactional processes in those calculi are associated with particular structures that record all process activities. Then, before granting a process the possibility of executing an action, the requested action is checked against the execution history to determine whether it will preserve consistency or not. Hence, different log definitions (in particular, the inference rules that check consistency) can provide different flavours of transactions.

Similarly, we leave out of the comparison the interesting work in (Bocchi and Tuosto 2010), where the basis are set for a theory of testing equivalence for distributed transactions in the presence of transactional attribute (inspired by the Java Transaction API). Transaction attributes discipline how services are executed with respect to the transactional scope of the invoking party. However, the calculus proposed in (Bocchi and Tuosto 2010) does define a notion of commit, but mostly focuses on compensation handling.

## Appendix A. Correctness and completeness of the implementation

Regarding the proofs of the theorems stating the correctness and completeness of the encoding of flat cJoin in Join, we report here the details omitted from the proof sketches in Section 5.

## A.1. *Correctness, part 1*

*Proof of Theorem 5.3.* The proof follows by case analysis on $P$. Note that $P$ cannot be of the form $[P' : Q']$ because $\nvdash [P' : Q'] : \square_1$. If $P$ has no local definitions (i.e., $P \not\equiv \textbf{def } D \textbf{ in } M$), then $P$ is either *abort*, the inert process 0, or the parallel composition of messages (containing only free names because there are no local definitions). In all three cases, $P' = P$. Last case is when $P$ contains local definitions, i.e., $P \equiv \textbf{def } D \textbf{ in } M$. We show that for any derivation $P \rightarrow^*_{cJ} P_1$ the following two conditions hold:

— $P_1 \equiv \textbf{def } D' \textbf{ in } M' \mid \Pi_{i \in 1..n} \mathcal{N}_i$, where $\mathcal{N}_i$ are cJoin transactions,
— $\exists Q_1$ s.t. $[\![P]\!] \rightarrow^*_J Q_1$ and $Q_1 \equiv (\textbf{def } [\![D']\!] \textbf{ in } [\![M']\!] \mid \Pi_{i \in 1..n} R_i) \mid \textbf{def } D_g \textbf{ in } 0$, where each $R_i$ is the standard Join negotiation associated to $\mathcal{N}_i$. Besides, $D_g$ are garbage definitions corresponding to instances of the commit protocol that have terminated.

The proof follows by induction on the length of the derivation $P \rightarrow^*_{cJ} P_1$. Base case ($m = 0$) follows immediately (since $P_1 = P$, it is enough to take $Q_1 = [\![P]\!]$). For the inductive step ($m = k + 1$), we consider reductions $P \rightarrow^k_{cJ} P'_1 \rightarrow_{cJ} P_1$. By inductive hypothesis on $P \rightarrow^k_{cJ} P'_1$ we know that:

1. $P'_1 \equiv \textbf{def } D'' \textbf{ in } M'' \mid \Pi_{i \in 1..n''} \mathcal{N}''_i$,
2. $\exists Q'_1$ s.t. $[\![P]\!] \rightarrow^*_J Q'_1$ and $Q'_1 \equiv (\textbf{def } [\![D'']\!] \textbf{ in } [\![M'']\!] \mid \Pi_{i \in 1..n''} R''_i) \mid \textbf{def } D''_g \textbf{ in } 0$, where each $R''_i$ is the standard Join negotiation associated to $\mathcal{N}''_i$.

We proceed by case analysis on the applied rule for $P'_1 \rightarrow_{cJ} P_1$. Interesting cases are those that terminate a cJoin transaction, i.e., when rule (COMMIT) or (ABORT) are used. Rule (COMMIT) can be only applied when $P'_1$ has a transaction $\mathcal{N}_1$ that does not contain *abort*, messages to local ports nor messages with local names as parameters. In this case, there exists $M''_1$ such that $\mathcal{N}''_1 \rightarrow_{cJ} M''_1$ and $P_1 \equiv \textbf{def } D'' \textbf{ in } M'' \mid M''_1 \mid \Pi_{i \in 2..n''} \mathcal{N}''_i$. By definition of a standard Join negotiation, $R''_1$ is such that all coordinators have been asked to commit and the set of continuations of such coordinators is $[\![M''_1]\!]$. We rely on a commit protocol that is ensured to terminate by releasing all continuations when all coordinators are asked to commit, then $R''_1 \rightarrow^* [\![M''_1]\!] \mid \textbf{def } D_{g1} \textbf{ in } 0$. Therefore, $Q'_1 \rightarrow^*$ $Q_1 \equiv (\textbf{def } [\![D'']\!] \textbf{ in } [\![M'']\!] \mid [\![M''_1]\!] \mid \Pi_{i \in 2..n''} R''_i) \mid \textbf{def } D''_g \textbf{ in } 0 \mid \textbf{def } D_{g1} \textbf{ in } 0$. Then, it is enough to take $D' = D''$, $M' = M'' \mid M''_1$, $n = n'' - 1$ with $\mathcal{N}'_i = \mathcal{N}''_{i+1}$ and $R'_i = R''_{i+1}$, and $D_g = D''_g \wedge D_{g1}$.

The case for (ABORT) follows analogously. $\square$

## A.2. *Completeness*

*Proof of Theorem 5.5.* We proceed by case analysis on the structure of $P$. Since $\vdash P : \square_1$, then $P \not\equiv [P' : Q']$. When $P$ has no local definitions, then it is the parallel composition of messages on free ports, the inert process 0 and *abort*. For any of these cases it holds that $[\![P]\!]$ does not have any definition, and therefore, $[\![P]\!]$ cannot reduce. The only possibility is $Q = norm(Q) = [\![P]\!]$, which trivially satisfies $\forall x : Q \downarrow_x \Rightarrow [\![P]\!] \downarrow_x$. If $P \equiv \textbf{def } D \textbf{ in } M$, then we show that the following three conditions hold:

1. $Q \equiv \textbf{def } [\![D']\!] \textbf{ in } [\![M'_1]\!] \mid \Pi_{i \in 1..u} R'_i \mid \Pi_{k \in 1..f} T'_k \mid \textbf{def } D_g \textbf{ in } 0$, where $R'_i$ are unfinished Join negotiations (i.e., some transactional thread has not finished), while $T'_k$ are finished negotiations, with $norm(\Pi_{k \in 1..f} T'_k) \equiv [\![M_2]\!] \mid \textbf{def } D_c \textbf{ in } 0$.

2. $P \to_{cJ}^* P' \equiv$ **def** $D'$ **in** $M_1' \mid M_2 \mid \Pi_{i \in 1..u} \mathcal{N}_i$ where $\mathcal{N}_i$ is a standard cJoin transaction corresponding to $R_i'$.

3. $norm(Q) \equiv$ **def** $[\![D']\!]$ **in** $[\![M_1'|M_2]\!] \mid norm(\Pi_{i \in 1..u} R_i') \mid$ **def** $D_g'$ **in** 0.

Above conditions can be shown by induction on the length of the derivation $[\![P]\!] \to_J^n Q$.

— **Base case** $Q = [\![P]\!]$. It is enough to take $P' = P$. Clearly, $P \to_{cJ}^* P' = P$. Since $\vdash P : \Box_1$, then $Q = [\![P]\!]$ has no coordinators (i.e., $n = 0$ and $f = 0$).

— **Inductive step** $[\![P]\!] \to_J^k Q'' \to_J Q$. By inductive hypothesis on $[\![P]\!] \to_J^k Q''$

1. $Q'' \equiv ;$ **def** $[\![D'']\!]$ **in** $[\![M_1'']\!] \mid \Pi_{i \in 1..u} R_i'' \mid \Pi_{k \in 1..F} T_k'' \mid$ **def** $D_g''$ **in** 0, where $R_i''$ are unfinished Join negotiations, $T_k''$ are finished negotiations and $norm(\Pi_{k \in 1..f} T_k'') \equiv [\![M_2'']\!]|$**def** $D_c''$ **in** 0.

2. $P \to_{cJ}^* P'' \equiv$ **def** $D''$ **in** $M_1'' \mid M_2'' \mid \Pi_{i \in 1..u} \mathcal{N}_i''$ where $\mathcal{N}_i''$ is the standard cJoin transaction corresponding to $R_i''$.

Then, the proof proceeds by case analysis of applied rule for reducing $Q'' \to_J Q$.

– for some $h \in 1..u''$, $R_h'' \to_J R_h'$. There are two different cases:

1. The applied rule corresponds to the commit protocol. Since, the protocol is confluent, then $norm(R_h') = norm(R_h'')$. Then, it is enough to take $P' = P''$, which satisfies all conditions.

2. The applied rule is not part of the commit protocol. Consequently, the applied rule is the encoding of some rule in $P$, and has the following shape $[\![x\langle \vec{u} \rangle \triangleright P_3]\!]$ or $[\![x\langle \vec{u} \rangle | x_1 \langle \vec{u}_1 \rangle \triangleright P_3]\!]$. We consider here the last case, which is the most interesting one. Hence, there exists a definition:

$$x\langle c_1, a_1, j_1 \rangle \mid y\langle c_2, a_2, j_2 \rangle \triangleright [\![P]\!]_{c_1, a_1, j_1} \mid j_1 \langle c_2, a_2 \rangle \mid j_2 \langle c_1, a_1 \rangle \mid c_2 \langle \rangle.$$

Moreover, $R_h''$ contains the messages for activating the rule. The application of the rule removes the consumed messages and activates the guarded process. The application of the rule will cause the two coordinators $j_1$ and $j_2$ to be joined to the same transaction. The effect of normalisation will depend on the structure of $P_3$:

a. $P_3 = y\langle \vec{v} \rangle$, s.t. $y$ is a message to a local port, then $R_h'$ contains $y\langle c_1, a_1, j_1 \rangle$ and the obtained transaction is unfinished. Clearly, this reduction corresponds to a reduction that merges two cJoin transactions.

b. If $P_3$ consists of a message to a merge port, then the proof is analogous to the previous case.

c. $[\![P_3]\!]c_1, a_1, j_1$ produces a commit vote, there are two cases: (i) if the vote is the last one, then, by normalising, $R_i$ commits. It is easy to notice that this case corresponds to the case in which all local names have been consumed, then there exist $P'$ s.t. $P'' \to_{cJ} P'$ by using commit; (ii) if some coordinators still wait the vote, then it is enough to take $P' = P''$.

d. $P_3 = abort$, then the encoding $[\![P_3]\!]c_1, a_1, j_1$ produces a commit vote to on the port $a_1$. The normalisation makes all coordinators in $R_h''$ to abort and to release the compensations. It is easy to notice that this corresponds

to $P'' \rightarrow^*_{cJ} P'$ by producing first the abort in the negotiation $h$ and then applying rule (ABORT), which releases all compensations.

1.  If the reduction is $\Pi_{k \in 1..f} T''_k \rightarrow_J \mathcal{R}$. Since, all $T''_k$ are finished negotiations and that normalisation procedure is confluent, then $norm(\Pi_{k \in 1..f} T''_k) = norm(\mathcal{R})$. Therefore, it is enough to take $P' = P''$.

2.  The applied rule is a definition in $[\![D'']\!]$:

    a.  the applied rule is part of the encoding of an ordinary definition:

        -   messages are in $[\![M''_1]\!]$, then immediate by reducing $P''$ by consuming messages in $M''_1$.

        -   if a message is in some $T''_k$. This is possible only if some coordinator has finished and released the continuation or the compensation, by correctness of the commit protocol, the message is in $[\![M''_2]\!]$, hence, it is possible to fire the corresponding rule in $P''$.

        Note that messages cannot be part of some $R_h$ because those transactions have not reached a decision, so global messages are kept by coordinators.

    b.  the applied rule is part of the encoding of a merge definition. This case is similar to the reduction internal to a negotiation and follows by analysing the pattern of the applied rule.

Finally, condition $\forall x : norm(Q) \downarrow_x \Rightarrow P' \downarrow_x$ immediately follows from conditions (2) and (3).

$\square$

## Appendix B. Formal definition of Coor

Before giving the full Join code for coordinators, we describe intuitively their behaviour with the transition state diagram in Figure 16. The initial state is called *state*. While in the initial state, a coordinator may accept requests for being joined (event *join*) with another participants. Any request is confirmed either with *okjoin* or *nojoin*. In both cases, the coordinator returns to the initial state. In the initial state, the coordinator can also receive the message to start the execution of the protocol, either with *cmt* (i.e., commit) or *abt* (i.e., abort). After receiving *cmt* the coordinator goes to the state *commit*. While in state *commit*, a coordinator behaves like in the original protocol, i.e., by notifying all known parties and by receiving commit confirmation until all parties commit. In such case, the coordinator reaches the state *finished*. Instead, if the coordinator receives the message *abt* when being in *state* or *commit*, it goes to state *abort*. While in *abort*, coordinators notify all known parties and discover the whole set of participants (analogously to commit). When all abort confirmations are received, the coordinator reaches the final state *finished*.

The Join code defining coordinators Coor is presented in Figure 17. Rule (0) fixes the initial state of the coordinator and is the only initially enabled rule of our encoding. This rule consumes the message $cmp\langle x \rangle$ and sets $x$ as the compensation to be activated on abort. The current state of the coordinator is represented with the message $state\langle \alpha, \beta \rangle$, where $\alpha$ is the compensation to be released on abort and $\beta$ is the list containing the
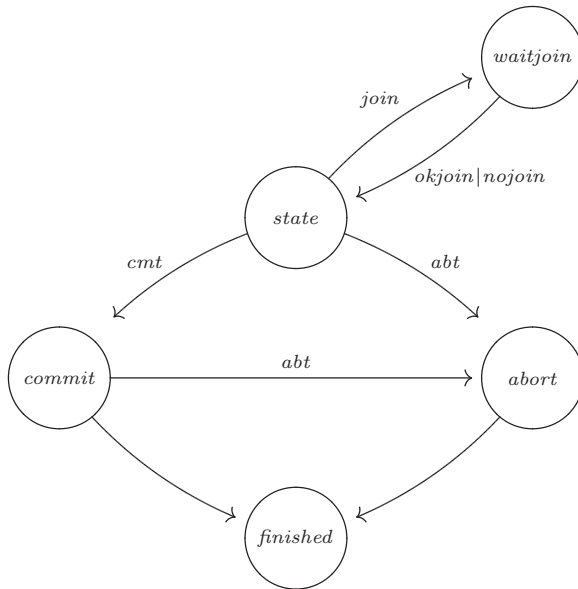
Fig. 16. States of coordinators.

channels corresponding to the coordinators of other parties in the same transaction (note that $\beta$ is initially empty). The following three rules (i.e., (1)–(3)) handle the joining of new parties in the transaction. When the coordinator is in the state $state\langle\alpha,\beta\rangle$ and receives a request $join\langle t, f\rangle$ for updating the state, it may accept the request (rule (1)) by passing to the state *waitjoin* and sends on $t$ the private ports on which it expects the update confirmation (i.e., message *okjoin*) or the cancellation (i.e., message *nojoin*). Rule (2) handles the reception of a join confirmation, which updates the set of known parties, while rule (3) deals with the cancellation. In both cases, the coordinator transits to the initial state (possibly updating it).

**Remark B.1.** For simplicity, we abstract away from this two-step communication in the presentation of Section 5 and we simply described join as a one-way message communication on port *join*.

Rule (4) starts the protocol with the commit vote, while rules (5)–(7) handle committing phase, and are analogous to the D2PC of (Bruni *et al.* 2002). There are two subtle differences: (i) channels *state* and *commit* have the extra parameter $\beta$, which is a list of the ports $abt_i$ of known participants to be used only if the state *abort* is reached; and (ii) coordinators goes to state *finished* after commit (rule (7)). Nevertheless, the behaviour for committing coordinators are as in the original proposal in (Bruni *et al.* 2002).

The behaviour for the aborting phase is given by rules (8)–(13). Rules (8) and (9) start the aborting phase when the coordinator receives a message on channel *abt* and it is either in the initial state (rule (8)) or in the commit phase (rule (9)). In both cases, the coordinator triggers a message $abort\langle\beta,\beta',\beta'',\alpha\rangle$, which carries the following values:

$$(0) \quad \mathsf{Coor} \equiv \qquad\qquad cmt\langle x\rangle \;\triangleright\; state\langle \{x\}, \emptyset\rangle$$

$$(1) \quad \wedge \qquad\qquad state\langle \alpha, \beta\rangle \,|\, join\langle t, f\rangle \;\triangleright\; t\langle okjoin, nojoin\rangle \,|\, waitjoin\langle \alpha, \beta\rangle$$

$$(2) \quad \wedge \qquad\qquad waitjoin\langle \alpha, \beta\rangle \,|\, okjoin\langle \beta'\rangle \;\triangleright\; state\langle \alpha, \beta \cup \beta'\rangle$$

$$(3) \quad \wedge \qquad\qquad waitjoin\langle \alpha, \beta\rangle \,|\, nojoin\langle\rangle \;\triangleright\; state\langle \alpha, \beta\rangle$$

$$(4) \quad \wedge \qquad\qquad state\langle \alpha, \beta\rangle \,|\, cmt\langle \ell, \kappa\rangle \;\triangleright\; commit\langle \ell \setminus \{lock\}, \ell, \{lock\}, \alpha, \kappa, \beta\rangle$$

$$(5) \quad \wedge \qquad commit\langle \{l\} \cup \ell, \ell', \ell'', \alpha, \kappa, \beta\rangle \;\triangleright\; commit\langle \ell, \ell', \ell'', \alpha, \kappa, \beta\rangle \,|\, l\langle \ell', lock, abt\rangle$$

$$(6) \quad \wedge \, commit\langle \ell, \ell', \ell'', \alpha, \kappa, \beta\rangle \,|\, lock\langle \ell''', l, a\rangle \;\triangleright\;$$
$$commit\langle \ell \cup (\ell''' \setminus \ell'), \ell' \cup \ell''', \ell'' \cup \{l\}, \alpha, \kappa, \beta \cup \{a\}\rangle$$

$$(7) \quad \wedge \qquad\qquad commit\langle \emptyset, \ell, \ell, \alpha, \kappa, \beta\rangle \;\triangleright\; \mathsf{release}\langle \kappa\rangle \,|\, finished\langle\rangle$$

$$(8) \quad \wedge \qquad\qquad state\langle \alpha, \beta\rangle \,|\, abt\langle \beta', a\rangle \;\triangleright\; abort\langle (\beta \cup \beta') \setminus \{abt\}, \beta \cup \beta', \{abt, a\}, \alpha\rangle$$

$$(9) \quad \wedge \quad commit\langle \emptyset, \ell', \ell'', \alpha, \kappa, \beta\rangle \,|\, abt\langle \beta', a\rangle \;\triangleright\;$$
$$abort\langle (\beta \cup \beta') \setminus \{abt\}, \beta \cup \beta', \{abt, a\}, \alpha\rangle \,|\, a\langle \beta, abt\rangle$$

$$(10) \quad \wedge \qquad\qquad abort\langle \{a\} \cup \beta, \beta', \beta'', \alpha\rangle \;\triangleright\; abort\langle \beta, \beta', \beta'', \alpha\rangle \,|\, a\langle \beta', abt\rangle$$

$$(11) \quad \wedge \qquad abort\langle \beta, \beta', \beta'', \alpha\rangle \,|\, lock\langle \ell''', l, a\rangle \;\triangleright\; abort\langle \beta \cup (\{a\} \setminus \beta'), \beta' \cup \{a\}, \beta'', \alpha\rangle$$

$$(12) \quad \wedge \qquad abort\langle \beta, \beta', \beta'', \alpha\rangle \,|\, abt\langle \beta', a\rangle \;\triangleright\;$$
$$abort\langle \beta \cup (\beta''' \setminus (\beta'' \cup \{a\})), \beta \cup \beta''', \beta'' \cup \{a\}, \alpha\rangle$$

$$(13) \quad \wedge \qquad\qquad abort\langle \emptyset, \beta, \beta, \alpha\rangle \;\triangleright\; \alpha\langle\rangle \,|\, finished\langle\rangle$$

$$(14) \quad \wedge \qquad\qquad finished\langle\rangle \,|\, cmt\langle \ell, cnt\rangle \;\triangleright\; finished\langle\rangle$$

$$(15) \quad \wedge \qquad\qquad finished\langle\rangle \,|\, lock\langle \ell, l, a\rangle \;\triangleright\; finished\langle\rangle$$

$$(16) \quad \wedge \qquad\qquad finished\langle\rangle \,|\, abt\langle \beta, a\rangle \;\triangleright\; finished\langle\rangle$$

$$(17) \quad \wedge \qquad\qquad finished\langle\rangle \,|\, join\langle t, f\rangle \;\triangleright\; f\langle\rangle \,|\, finished\langle\rangle$$

$$(18) \quad \wedge \qquad commit\langle \ell, \ell', \ell'', \alpha, \kappa, \beta\rangle \,|\, join\langle t, f\rangle \;\triangleright\; f\langle\rangle \,|\, commit\langle \ell, \ell', \ell'', \alpha, \kappa, \beta\rangle$$

$$(19) \quad \wedge \qquad\qquad abort\langle \beta, \beta', \beta'', \alpha\rangle \,|\, join\langle t, f\rangle \;\triangleright\; f\langle\rangle \,|\, abort\langle \beta, \beta', \beta'', \alpha\rangle$$

Fig. 17. Join code of coordinators.

— $\beta$ records the set of *abt* ports of known participants that must still be contacted (analogous to $\ell$);
— $\beta'$ stores the list of ports *abt*$_i$ of known participants involved in the same transaction, which is typically augmented during the D2PC with the sets sent by other participants (analogous to $\ell'$);

— $\beta''$ records the parties who have already sent their consensus for abort (analogous to $\ell''$);

— $\alpha$ store the messages to be released when aborting, i.e., the activation of the compensation.

Note that the behaviour for the aborting phase (rules (10)–(13)) is analogous to the committing phase, and it can be described as follows:

1. **First phase**. The participant sends the abort vote to every known thread in $\beta$ (rule (10)). The message contains the list $\beta'$ of all known participants, and the sender identification *abt*.

2. **Second phase**. The participant collects the messages sent by other parties and updates its own synchronisation set (rule (11) and (12)). A request will be also sent to the new items in the synchronisation set (by repeating the first phase for them).

3. When the set of aborting parties is transitively closed, the protocol terminates locally and the coordinator transits to the state *finished* and releases the compensation $\alpha$ (rule (13)).

Rules (14)–(16) are for collecting garbage, and state that messages received when the protocol has finished are ignored. Moreover, rules (17)–(19) state that the state of a coordinator cannot be updated when the protocol has begun.

## References

Anderson, B. and Shasha, D. (1992) Persistent linda: Linda + transactions + query processing. In: *Research Directions in High-Level Parallel Programming Languages*, Springer Verlag 93–109.

Benton, N., Cardelli, L. and Fournet, C. (2002) Modern concurrency abstractions for C♯. In: Proceedings of ECOOP 2002. *Springer Verlag Lecture Notes in Computer Science* **2374** 415–440.

Bernstein, P., Hadzilacos, V. and Goodman, N. (1987) *Concurrency, Control and Recovery in Database Systems*, Addison-Wesley Longman.

Berry, G. (1993) Preemption in concurrent systems. In: Proceedings of FSTTCS'93. *Springer Verlag Lecture Notes in Computer Science* **761** 72–93.

Berry, G. and Boudol, G. (1992) The chemical abstract machine. *Theoretical Computer Science* **96** (1) 217–248.

Bocchi, L., Laneve, C. and Zavattaro, G. (2003) A calculus for long-running transactions. In: Proceedings of FMOODS'03. *Springer Verlag Lecture Notes in Computer Science* **2884** 124–138.

Bocchi, L. and Tuosto, E. (2010) Testing attribute-based transactions in SOC. In: Proceedings of FMOODS/FORTE 2010. *Springer Verlag Lecture Notes in Computer Science* **6117** 87–94.

Boreale, M., Bruni, R., De Nicola, R. and Loreti, M. (2008) Sessions and pipelines for structured service programming. In: Barthe, G. and de Boer, F. S. (eds.) Proceedings of FMOODS'08. *Springer Verlag Lecture Notes in Computer Science* **5051** 19–38.

BPEL (2003) BPEL Specification. version 1.1. Available at `http://www.ibm.com/developerworks/library/ws-bpel`.

BPMN (2010) Business process modelling notation (BPMN). Available at `http://www.bpmi.org`.

Bruni, R., Kersten, A., Lanese, I. and Spagnolo, G. (2012) A new strategy for distributed compensations with interruption in long-running transactions. In: Proceedings of WADT 2010. *Springer Verlag Lecture Notes in Computer Science* **7137** 42–60.

Bruni, R., Laneve, C. and Montanari, U. (2002) Orchestrating transactions in join calculus. In: Proceedings of CONCUR 2002. *Springer Verlag Lecture Notes in Computer Science* **2421** 321–336.

Bruni, R., Melgratti, H. and Montanari, U. (2003) Flat committed join in join. In: Proceedings of CoMeta 2003. *Elsevier Science Electronic Notes in Theoretical Computer Science* **104** 39–59.

Bruni, R., Melgratti, H. and Montanari, U. (2004) Nested commits for mobile calculi: Extending Join. In: *Proceedings of the 3rd IFIP-TCS 2004*, Kluwer Academic Publishers 569–582.

Bruni, R., Melgratti, H. and Montanari, U. (2005) Theoretical foundations for compensations in flow composition languages. In: *Proceedings of POPL 2005*, ACM Press 209–220.

Bruni, R. and Montanari, U. (2004) Concurrent models for Linda with transactions. *Mathematical Structures in Computer Science* **14** (3) 421–468.

Busi, N. and Zavattaro, G. (2002) On the serializability of transactions in shared dataspaces with temporary data. In: *Proceedings of SAC 2002*, ACM Press 359–366.

Butler, M., Bruni, R., Ferreira, C., Hoare, T., Melgratti, H. and Montanari, U. (2005a) Comparing two approaches to compensable flow composition. In: Proceedings of CONCUR 2005. *Springer Verlag Lecture Notes in Computer Science* **3653** 383–397.

Butler, M., Chessell, M., Ferreira, C., Griffin, C., Henderson, P. and Vines, D. (2002) Extending the concept of transaction compensation. *IBM Systems Journal* **41** (4) 743–758.

Butler, M. and Ferreira, C. (2004) An operational semantics for StAC, a language for modelling long-running business transactions. In: Proceedings of Coordination 2004. *Springer Verlag Lecture Notes in Computer Science* **2949** 87–104.

Butler, M., Hoare, T. and Ferreira, C. (2005b). A trace semantics for long-running transactions. In: Proceedings of 25 Years of CSP. *Springer Verlag Lecture Notes in Computer Science* **3525** 133–150.

Caires, L., Ferreira, C. and Vieira, H.T. (2009) A process calculus analysis of compensations. In: Kaklamanis, C. and Nielson, F. (eds.) Proceedings of TGC'08. *Springer Verlag Lecture Notes in Computer Science* **5474** 87–103.

Chothia, T. and Duggan, D. (2004) Abstractions for fault-tolerant global computing. *Theoretical Computer Science* **322** (3) 567–613.

Conchon, S. and Le Fessant, F. (1999) Jocaml: Mobile agents for Objective-Caml. In: *Proceedings of ASA/ MA'99*. IEEE Computer Society 22–29.

Danos, V. and Krivine, J. (2004) Reversible communicating systems. In: Proceedings of CONCUR 2004. *Springer Verlag Lecture Notes in Computer Science* **3170** 293–307.

de Vries, E., Koutavas, V. and Hennessy, M. (2010) Communicating transactions - (extended abstract). In: Gastin, P. and Laroussinie, F, (eds.) Proceedings of CONCUR'10. *Springer Verlag Lecture Notes in Computer Science* **6269** 569–583.

Eisentraut, C. and Spieler, D. (2009) Fault, compensation and termination in ws-bpel 2.0 - a comparative analysis. In: Bruni, R. and Wolf, K. (eds.) Proceedings of WS-FM'08. *Springer Verlag Lecture Notes in Computer Science* **5387** 107–126.

Elmagarmid, A., Leu, Y., Litwin, W. and Rusinkiewicz, M. (1990) A multidatabase transaction model for interbase. In: *Proceedings of VLDB'90*. Morgan Kaufmann 507–518.

Eswaran, K., Gray, J., Lorie, R. and Traiger, I. (1976) The notions of consistency and predicate locks in a database system. *Communications of the ACM* **19** (11) 624–633.

Fekete, A., Lynch, N., Merritt, M. and Weihl, W. (1994) *Atomic Transactions*, Morgan Kaufmann Publishers.

Fournet, C. and Gonthier, G. (1996) The reflexive chemical abstract machine and the Join calculus. In: *Proceedings of POPL'96*. ACM Press 372–385.

Fournet, C., Gonthier, G., Lévy, J.-J., Maranget, L. and Rémy, D. (1996) A calculus of mobile agents. In: Proceedings of CONCUR'96. *Springer Verlag Lecture Notes in Computer Science* **1119** 406–421.

Garcia-Molina, H. and Salem, K. (1987) Sagas. In: *Proceedings of the ACM Special Interest Group on Management of Data Annual Conference*. ACM Press 249–259.

Georgakopoulos, D., Hornick, M. and Sheth, A. (1995) An overview of workflow management: From process modeling to workflow automation infrastructure. *Distributed and Parallel Databases* **3** (2) 119–153.

Gray, J. and Reuter, A. (1993) *Transaction Processing: Concepts and Techniques*, Morgan Kaufmann.

Hutchinson, N., Kaiser, G. and Pu, C. (1988) Split-transactions for open-ended activities. In: *Proceedings of VLDB'88*, Morgan Kaufmann 26–37.

Jagannathan, S. and Vitek, J. (2004) Optimistic concurrency semantics for transactions in coordination languages. In: Proceedings of COORDINATION 2004, *Springer Verlag Lecture Notes in Computer Science* **2949** 183–198.

Kaiser, G. and Pu, C. (1992) Dynamic restructuring of transactions. In: *Database Transaction Models for Advanced Applications*, Morgan Kaufmann 265–295.

Kochut, K., Miller, J., Sheth, A. and Wang, X. (1996) Corba-based run-time architectures for workflow management systems. *Journal of Database Management, Special Issue on Multidatabases* **7** (1) 16–27.

Kohler, W. (1981) A survey of techniques for synchronization and recovery in decentralized computer systems. *ACM Computing Surveys* **13** (2) 149–183.

Krishnakumar, N. and Sheth, A. (1995) Managing heterogeneous multi-system tasks to support enterprise-wide operations. *Distributed and Parallel Databases* **3** (2) 155–186.

Lanese, I., Mezzina, C. and Stefani, S. (2010a) Reversing higher-order pi. In: Proceedings of CONCUR'10. *Springer Verlag Lecture Notes in Computer Science* **6269** 478–493.

Lanese, I., Vaz, C. and Ferreira, C. (2010b) On the expressive power of primitives for compensation handling. In: Proceedings of ESOP'10. *Springer Verlag Lecture Notes in Computer Science* **6012** 366–386.

Laneve, C. and Zavattaro, G. (2005) Foundations of web transactions. In: Proceedings of FOSSACS 2005. *Springer Verlag Lecture Notes in Computer Science* **3441** 282–298.

Leymann, F. (2001) WSFL Specification. version 1.0. Available at http://www-306.ibm.com/software/solutions/webservices/pdf/WSFL.pdf.

Lomet, D. (1992) MLR: A recovery method for multi-level systems. In: *Proceedings of the 1992 ACM SIGMOD International Conference on Management of Data*. ACM Press 185–194.

Lucchi, R. and Mazzara, M. (2004) A framework for generic error handling in business processes. In: Proceedings of WS-FM 2004. *Electronic Notes in Theoretical Computer Science* **105** 133–145. (Elsevier Science.)

Melgratti, H. (2005) *Models and Languages for Global Computing Transactions*, Ph.D. thesis, Computer Science Department, University of Pisa.

Milner, R. (1980) *A Calculus of Communicating Systems Springer Verlag Lecture Notes in Computer Science* **92**.

Milner, R., Parrow, J. and Walker, J. (1992) A calculus of mobile processes, I and II. *Information and Computation* **100** (1) 1–77.

Moss, J. (1981) *Nested Transactions: An Approach to Reliable Distributed Computing*, Ph.D. thesis, Dept. of Electrical Engineering and Computer Science, MIT.

Reuter, A. and Wächter, H. (1992) The contract model. *Transaction Models for Advanced Applications*, Morgan Kaufmann 219–263.

Ripon, S. (2008) *Extending and Relating Semantic Models of Compensating CSP*, Ph.D. thesis, School of Electonics and Computer Science, University of Southampton.

Rusinkiewicz, M. and Sheth, A. (1995) Specification and execution of transactional workflows. In: *Modern Database Systems: The Object Model, Interoperability, and Beyond*. ACM Press and Addison-Wesley 592–620.

Schek, H.-J. and Weikum, G. (1992) Concepts and applications of multilevel transactions and open nested transactions. In: *Database Transaction Models for Advanced Applications*. Morgan Kaufmann 515–553.

Thatte, S. (2001) XLANG: Web Services for Business Process Design. Available at `http://www.gotdotnet.com/team/xml_wsspecs/xlang-c/default.htm`.

Vaz, C., Ferreira, C. and Ravara, A. (2008) Dynamic recovering of long running transactions. In: Proceedings of TGC 2008. *Springer Verlag Lecture Notes in Computer Science* **5474** 201–215.

Weikum, G. (1991) Principles and realization strategies of multilevel transaction management. *ACM Transactions on Database Systems* **16** (1) 132–180.

WSCDL (2004) Web Services Choreography Description Language. Version 1.0. Available at `http://www.w3.org/TR/2004/WD-ws-cdl-10-20040427/`.

WSCI (2002) WSCI Specification. version 1.0. Available at `http://www.w3.org/TR/wsci/`.