

NESTED COMMITS FOR MOBILE CALCULI: EXTENDING JOIN *

Roberto Bruni, Hernán Melgratti, Ugo Montanari
Dipartimento di Informatica, Università di Pisa, Italia.
{bruni, melgratt, ugo}@di.unipi.it

Abstract In global computing applications the availability of a mechanism for some form of committed choice can be useful, and sometimes necessary. It can conveniently handle, e.g., distributed agreements and negotiations with nested choice points. We propose a linguistic extension of the Join calculus for programming nested commits, called *Committed Join* (cJoin). It provides primitives for explicit abort, programmable compensations and interactions between negotiations. We give the operational semantics of cJoin in the *reflexive CHAM* style. Then we discuss its expressiveness on the basis of a few examples and encodings. Finally, we provide a big-step semantics for cJoin processes that can be typed as *shallow* and we show that shallow processes are serializable.

1. Introduction

In recent years, wide area network computing, web programming, and, more generally, *global computing* (GC) are attracting the interest of many researchers in an attempt of laying the foundations for largely distributed applications. Such applications often require a coordination layer to orchestrate their components, which are designed and implemented separately, run on different platforms and communicate asynchronously. Often, the components must agree on the activities they are carrying on (e.g. in terms of transactions, like in [16, 7]) by committing the results of long distributed decision processes as soon as the participants reach partial agreements. Applications can handle these situations in an ad hoc manner or they can rely on a fixed set of coordination primitives. In this work we are interested on studying suitable primitives for describing distributed commits in GC applications. Note that we use the term “*commit*” (also *contract* or *negotiation*) instead of “*transaction*” to emphasize the coordination aspects, which are orthogonal to ACID database transactions. For instance, in

* Research supported by the MSR Cambridge Project NAPI, by the FET-GC Project IST-2001-32747 AGILE, by the MIUR Projects COFIN CoMETA and IS-MANET, and by the MURST-CNR 1999 Project.

the case of web services, the orchestration layer should provide the primitives for specifying transactional services and the valid interactions between them. Nevertheless, any service should be responsible for maintaining the consistency on their local data (i.e., assuring ACID properties on them).

Process description languages (PDLs) are mathematical models of computation designed for isolating and studying phenomena that occur in concurrent languages. In the spirit of PDL, it would be desirable to extend well-known calculi with primitives for distributed nested commits. Two key operations are the “abort with compensation” (e.g., to stop a negotiation when some participants withdraw their interest in carrying out the contract) and the “commit” (to store a partial agreement before moving to the next phase of a long negotiation). In this paper we introduce *Committed Join* (cJoin) as an extension of the Join calculus [12]. The features of cJoin are compared against two other paradigms with commit, namely AKL [13] and *Zero-Safe* nets [8].

The design of cJoin has been inspired by the requirements (i)–(vi) below. Contracts are decision processes distributed on several nodes, each with the possibility of consulting both local and global resources and generating local sub-contracts (e.g. modeling decisions internal to an organization). However: (i) each internal sub-decision should be stored locally and not made public before a common agreement is achieved; (ii) global resources might be made available to partners upon commits, marking the conclusion of some contract; (iii) decision processes can be aborted, in which case all participants should be informed and suitable compensation procedures activated (e.g., upon abort, new attempts for finding an agreement can be initiated) ; (iv) divergence is possible, but well designed GC applications should guarantee that each contract eventually leads to an abort or a commit; (v) when two processes involved in separate negotiations exchange some information, then their contracts should be merged into a unique one; (vi) it should be possible to have nested contracts. Though an internal abort can be compensated in such a way that the main contract can still be successfully completed, a failure of the main contract should cause the abort of all ongoing internal activities.

We define the small-step operational semantics of cJoin in the *reflexive* CHAM style [12]. We also give a big-step semantics for the sub-calculus of *shallow* processes and we show that shallow processes are serializable by proving a correspondence between their CHAM and big-step semantics. *Serializability* ensures the correctness of reasoning at different levels of abstractions when transactions become atomic transitions at the abstract level.

Synopsis. In § 2 we recall the principles of the CHAM and the syntax and semantics of Join. Committed Join is introduced in § 3. In § 4 we illustrate the main features of cJoin by showing a simple application for booking trips and the encoding of AKL. The implementation of ZS nets is presented in § 5. Finally, we study *serializability* in § 6.

2. Background: CHAM and Join

The Chemical Abstract Machine. In CHAM [3] states s (called *solutions*) are finite multisets of terms m (called *molecules*), and computations are multiset rewrites. Multisets are denoted by m_1, \dots, m_n and abbreviated with $\otimes_i m_i$. Solutions can be structured in a hierarchical way by using the operator *membrane* $\{\cdot\}$ to group a solution s into a molecule $\{\{s\}\}$. In [3] 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* (or *structural*) rules \rightleftharpoons representing syntactical rearrangements of molecules in a solution, and *reaction* rules \rightarrow . Structural rules are reversible: a solution obtained by applying a cooling rule can be heated back to the original state, and vice versa. Instead, reaction rules cannot be undone.

The laws governing CHAM computations state that whenever an instance of a left-hand-side of a rule is found into a solution, it can be replaced by the corresponding instance of the right-hand-side. Chemical rules $S_1 \rightarrow S_2$ have no premises and are purely local, specifying only the part of the solution that actually changes. Consequently, they can be applied in every larger solution $S, S_1 \rightarrow S, S_2$ (*chemical law*) and also in grouped sub-solutions, $\{\{S, S_1\}\} \rightarrow \{\{S, S_2\}\}$. In particular, they can be nested at any level of hierarchical solutions. Note that, since solutions are multisets, rules can be applied concurrently.

The Join calculus. The Join calculus [12] is a well-known PDL with asynchronous name-passing communication. It has the same expressive power as the asynchronous π -calculus and it has distributed running implementations, e.g. *Jocaml* [10] and *Polyphonic C[#]* [1]. Join relies on an infinite set of names x, y, u, v, \dots . Name tuples are written \vec{u} . Join processes, definitions and patterns are in Figure 1.a. A *process* is either the inert process 0 , the asynchronous emission $x\langle\vec{y}\rangle$ of message \vec{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 defined as usual.

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 correspond to terms of the Join calculus denoting processes or definitions. The chemical rules are shown in Figure 1.b. Rule STR-NULL states that 0 can be added or removed from any solution. Rules STR-JOIN and STR-AND stand for the associativity and commutativity of $|$ and \wedge , because $-, -$ is such. STR-DEF denotes the activation

<p>(PROC) $P, Q ::= 0 \mid x\langle \vec{y} \rangle \mid \mathbf{def} D \mathbf{in} P \mid P Q$</p> <p>(DEF) $D, E ::= J \triangleright P \mid D \wedge E$</p> <p>(PAT) $J, K ::= x\langle \vec{y} \rangle \mid J K$</p> <p style="text-align: center;">a. Syntax.</p>	<p>(STR-NULL) $0 \Leftarrow$</p> <p>(STR-JOIN) $P \mid Q \Leftarrow P, Q$</p> <p>(STR-AND) $D \wedge E \Leftarrow D, E$</p> <p>(STR-DEF) $\mathbf{def} D \mathbf{in} P \Leftarrow D\sigma_{dn(D)}, P\sigma_{dn(D)}$ <i>(range($\sigma_{dn(D)}$) globally fresh)</i></p> <p>(RED) $J \triangleright P, J\sigma \rightarrow J \triangleright P, P\sigma$</p> <p style="text-align: center;">b. Semantics.</p>
---	--

Figure 1. Join Calculus.

<p>$M, N ::= 0 \mid x\langle \vec{y} \rangle \mid M N$</p> <p>$P, Q ::= M \mid \mathit{abort} \mid \mathbf{def} D \mathbf{in} P \mid [P : Q] \mid P Q$</p> <p>$D, E ::= J \triangleright P \mid J \blacktriangleright P \mid D \wedge E$</p> <p>$J, K ::= x\langle \vec{y} \rangle \mid J K$</p>	Figure 2 Syntax of cJoin.
--	---------------------------

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 $x_1 \dots x_n$ by names $y_1 \dots y_n$ as $\sigma = \{y_1 \dots y_n / x_1 \dots x_n\}$, with $dom(\sigma) = \{x_1, \dots, x_n\}$ and $range(\sigma) = \{y_1, \dots, y_n\}$. We indicate with σ_N an injective substitution σ such that $dom(\sigma) = N$. We require names to be globally fresh, i.e. fresh w.r.t. the implicit context in which the rule is applied.

Consider, for instance, $s = \{z\langle x, z \rangle, \mathbf{def} x\langle y \rangle \triangleright z\langle y, x \rangle \mathbf{in} x\langle a \rangle\}$, whose second molecule contains a definition of a local port x different from the homonym free port in the first molecule. When STR-DEF is applied, the local definition of x is renamed by using a fresh name, obtaining, for instance, the solution $s' = \{z\langle x, z \rangle, x_1\langle y \rangle \triangleright z\langle y, x_1 \rangle, x_1\langle a \rangle\}$.

Finally, RED describes the use of an active reaction rule ($J \triangleright P$) to consume messages forming an instance of J (for a suitable substitution σ , with $dom(\sigma) = rn(J)$), and produce a new instance $P\sigma$ of its guarded process P . By applying RED to s' for $\sigma = \{a / y\}$, we get $s' \rightarrow \{z\langle x, z \rangle, x_1\langle y \rangle \triangleright z\langle y, x_1 \rangle, z\langle a, x_1 \rangle\}$. Note that the local port x_1 has been extruded on the free channel z .

3. Committed Join

Syntax. We extend the syntax of JOIN as in Figure 2. A negotiation is represented by $[P : Q]$, where P is the normal activity and Q is its compensation. The normal activity P is intended to execute in isolation until reaching either a commit or an abort decision. If P commits, the obtained result is delivered to the outside of the negotiation. Instead, Q is activated when P aborts. The abort decision is signaled with the special process *abort*.

A new kind of definitions $J \blacktriangleright P$, called *merge definitions*, is introduced to describe the interactions among negotiations. Merge definitions allow the

$$\begin{aligned} dn_o(D \wedge E) &= dn_o(D) \cup dn_o(E) & dn_o(J \triangleright P) &= dn(J) & dn_o(J \blacktriangleright P) &= \emptyset \\ dn_m(D \wedge E) &= dn_m(D) \cup dn_m(E) & dn_m(J \triangleright P) &= \emptyset & dn_m(J \blacktriangleright P) &= dn(J) \end{aligned}$$

Figure 3. Ordinary and merge names.

$$\begin{array}{l} \text{(STR-NULL)} \quad 0 \rightleftharpoons \\ \text{(STR-JOIN)} \quad P \mid Q \rightleftharpoons P, Q \\ \text{(STR-AND)} \quad D \wedge E \rightleftharpoons D, E \\ \text{(STR-DEF)} \quad \mathbf{def} D \mathbf{in} P \rightleftharpoons D \sigma_{dn(D)}, P \sigma_{dn(D)} \text{ (range}(\sigma_{dn(D)}) \text{ globally fresh)} \\ \text{(RED)} \quad J \triangleright P, J \sigma \rightarrow J \triangleright P, P \sigma \\ \\ \text{(STR-CONT)} \quad [P : Q] \rightleftharpoons \{\!\{P, \perp Q\!\}\!\} \\ \text{(COMMIT)} \quad \{\!\{M \mid \mathbf{def} D \mathbf{in} 0, \perp Q\!\}\!\} \rightarrow M \\ \text{(ABORT)} \quad \{\!\{abort \mid P, \perp Q\!\}\!\} \rightarrow Q \\ \text{(MERGE)} \quad \Pi_j J_j \blacktriangleright P, \otimes_i \{J_i \sigma, S_i, \perp Q_i\} \rightarrow \Pi_j J_j \blacktriangleright P, \{\!\{\otimes_i S_i, P \sigma, \perp \Pi_i Q_i\!\}\!\} \end{array}$$

Figure 4. Operational semantics of cJoin.

consumption of messages produced in the scope of different contracts by joining all participants in a unique larger negotiation. Moreover, usual definitions can be used to create negotiations dynamically. For instance, by firing $J \triangleright [P : Q]$ a new instance of the negotiation P with compensation Q is activated.

For convenience we introduce the syntactical category M of processes without definitions, i.e. a parallel composition of messages.

The definition of fn is extended with $fn(abort) = \emptyset$, $fn([P : Q]) = fn(P) \cup fn(Q)$ and $fn(J \blacktriangleright P) = dn(J) \cup (fn(P) \setminus rn(J))$. For a definition D , we redefine $dn(D) = dn_o(D) \cup dn_m(D)$, where dn_o denotes the *defined ordinary* names and dn_m the *defined merge* names (Figure 3). We assume $dn_o(D) \cap dn_m(D) = \emptyset$ for every definition D .

Operational Semantics. The operational semantics of cJoin is defined in the reflexive CHAM style. Molecules m and solutions S are defined below.

$$m ::= P \mid D \mid \perp P \mid \{\!\{S\}\!\} \quad S ::= m \mid m, S$$

As in ordinary Join, processes and definitions are molecules. Additionally, a molecule $\perp P$ denotes a compensation that is frozen inside a solution. The chemical rules are in Figure 4. The first five rules are the ordinary ones for Join. Rule STR-CONT describes how a negotiation corresponds to a sub-solution of two molecules: the process P and its compensation Q , which is frozen (because the operator \perp forbids the enclosed process to react).

COMMIT can be executed only when all participants have done their tasks reaching a (local) state that does not contain locally defined names. This way, a commit means clean termination where all names denoting internal states of contracts have been consumed. Note that all definitions belonging to a contract are discarded at commit time because the messages that are being released do not contain those names (we recall that local names cannot be extruded). Similarly, the compensation is discarded at commit. Moreover, a negotiation cannot

commit when *abort* is within the solution because *abort* is not a message. The abort is handled by rule ABORT, which activates the compensation procedure while discarding all terms in the solution. Compensations are not predefined to execute atomically, but they can be explicitly programmed as negotiations.

Interactions among contracts are specified by rule MERGE, which consumes messages $J_i\sigma$ from different contracts and creates a new larger negotiation by combining the existing contracts together with the new instance $P\sigma$, where $dom(\sigma) = rn(\Pi_j J_j)$. The compensation for the joint negotiation is the parallel composition of all the original compensations. When merging negotiations, clashes of locally defined names should be avoided by imposing the side condition $dn(S_i) \cap (dn(S_j) \cup fn(S_j)) = \emptyset$ for $i \neq j$. However, if we are guaranteed that STR-DEF generates globally fresh names (and not just locally fresh names) then this side condition can be safely omitted, as it is trivially satisfied.

PROPOSITION 1 *cJoin* is a conservative extension of *Join*.

Discussion. Sibling contracts can be merged only by using merge definitions introduced by their parent. In practice, it might be useful to apply a merge definition provided by any ancestor. To this aim, the rule STR-MOVE below might be added, so that merge definitions could float across contract boundaries.

$$\text{STR-MOVE} \quad J \blacktriangleright P, \{S\} \rightleftharpoons \{J \blacktriangleright P, S\}$$

Regarding deadlocks, note that stall negotiations are not discarded. For instance, the process $[\mathbf{def} \ x \langle \rangle | y \langle \rangle \triangleright 0 \ \mathbf{in} \ x \langle \rangle : Q]$ cannot compute. Neither it can commit, because there is a message on the local port x . In this situation the contract is blocked and should be aborted. Some of these situations can be recognized and handled locally to promote the abort (i.e., when no local rules can be applied). These situations can be represented by ad hoc rules or by a general rule to generate nondeterministically the abort (situations that real implementations typically handle with timeouts). Nevertheless, we cannot expect to axiomatize stall situations because it would mean to write axioms recognizing non-termination, which is an undecidable problem. On the other hand, we do not want to limit the expressiveness of the language.

With respect to the requirements discussed in the Introduction, we have that, membranes are exploited to define the boundaries of negotiations. Process like $[P_1 \mid [P_2 : Q_2] \mid [P_3 : Q_3] : Q_1]$ straightforwardly model sub-negotiations. The decisions taken internally by P_2 can influence P_3 only if some merge definition is available at the level of P_1 . In absence of merge definitions, global and local resources are kept separate in each sub-negotiation. The commit of P_2 can only happen when the internal state contains only global resources. At commit time, the result of the negotiation is made available to P_1 . An abort generated in P_2 activates the compensation Q_2 at the level of P_1 , neither forcing the abort of

$$\begin{aligned}
\mathbf{H} &\equiv \mathbf{def} \quad \mathit{WaitBooking}\langle \rangle \triangleright [\mathbf{def} && \mathit{request}\langle o \rangle \triangleright o\langle \$ \rangle \mid \mathit{price}\langle \$ \rangle \\
&& \wedge \mathit{price}\langle \$ \rangle \mid \mathit{confirm}\langle v \rangle \triangleright \mathit{BookedRoom}\langle v \rangle \\
&& \wedge \mathit{price}\langle \$ \rangle \triangleright \mathit{abort} \\
&& \mathbf{in} \quad \mathit{offerRoom}\langle \mathit{request}, \mathit{confirm} \rangle : Q] \\
&\wedge \mathit{BookedRoom}\langle v \rangle \triangleright R \\
&\mathbf{in} \quad \mathit{WaitBooking}\langle \rangle \mid \dots \\
\mathbf{C} &\equiv \mathbf{def} \quad \mathit{BookHotel}\langle \rangle \triangleright [\mathbf{def} \mathit{HotelMsg}\langle r, c \rangle \triangleright \\
&& \mathbf{def} \mathit{offer}\langle \$ \rangle \triangleright c\langle \mathit{visa} \rangle \mid \mathit{hotelOK}\langle hConf \rangle \\
&& \wedge \mathit{offer}\langle \$ \rangle \triangleright \mathit{abort} \\
&& \wedge hConf\langle \rangle \triangleright \mathit{HotelFound}\langle \rangle \\
&& \mathbf{in} \quad r\langle \mathit{offer} \rangle \quad \mathbf{in} \quad \mathit{searchRoom}\langle \mathit{HotelMsg} \rangle : Q'] \\
&\wedge \mathit{BookFlight}\langle \rangle \triangleright [\mathbf{def} \mathit{FlightMsg}\langle r, c \rangle \triangleright \\
&& \mathbf{def} \mathit{offer}\langle \$ \rangle \triangleright c\langle \mathit{visa} \rangle \mid \mathit{flightOK}\langle fConf \rangle \\
&& \wedge \mathit{offer}\langle \$ \rangle \triangleright \mathit{abort} \\
&& \wedge fConf\langle \rangle \triangleright \mathit{FlightFound}\langle \rangle \\
&& \mathbf{in} \quad r\langle \mathit{offer} \rangle \quad \mathbf{in} \quad \mathit{searchFlight}\langle \mathit{FlightMsg} \rangle : Q''] \\
&\wedge \mathit{flightOK}\langle fc \rangle \mid \mathit{hotelOK}\langle rc \rangle \blacktriangleright fc\langle \rangle \mid rc\langle \rangle \\
&\mathbf{in} \quad \mathit{BookHotel}\langle \rangle \mid \mathit{BookFlight}\langle \rangle \mid \dots \\
\mathbf{Trip} &\equiv \mathbf{def} \quad \mathit{searchRoom}\langle hm \rangle \mid \mathit{offerRoom}\langle r, c \rangle \blacktriangleright hm\langle r, c \rangle \\
&\wedge \mathit{searchFlight}\langle fm \rangle \mid \mathit{offerFlight}\langle r, c \rangle \blacktriangleright fm\langle r, c \rangle \quad \mathbf{in} \quad \mathbf{H} \mid \mathbf{A} \mid \mathbf{C}
\end{aligned}$$

Figure 5. Trip booking example.

any other sub-negotiations (P_3) nor the abort of the main contract (P_1). Note that if $[P_2 : Q_2]$ was the result of the merging of several negotiations, then Q_2 is the union of all the compensations of the participants.

An important restriction is that local resources can neither cross negotiations boundaries nor be extruded to siblings negotiations. The only way to exchange information between siblings negotiations is by merging all participants into a unique negotiation that must then commit, or abort, or diverge as such.

4. Examples

Trip booking. Figure 5 shows the encoding of the application Trip that allows a user to book flights and accommodations. Trip is defined in term of three components: the hotel H, the airline A and the customer C. The component H is a process that activates (by firing the definition for $\mathit{WaitBooking}$) a negotiation to serve customer requests. Such negotiation starts by publishing on the merge port $\mathit{offerRoom}$ (defined in Trip) the names of the services a client should use to reserve a room: $\mathit{request}$ to ask for a quote; and $\mathit{confirm}$ to accept an offer. The component A (omitted in Figure 5) is defined analogously, but it publishes services on port $\mathit{offerFlight}$ instead of $\mathit{offerRoom}$.

The component C defines two rules for creating negotiations: one for booking rooms and the other for buying flight tickets. Both contracts are quite similar. In particular, the negotiation for booking a room starts by sending a message to the merge port $\mathit{searchRoom}$ (defined in Trip) to obtain the names for interacting

with a hotel. The first merge rule in Trip will associate an offer from a hotel with a request from a client by sending the names r and c to the corresponding port hm . Once received r and c on *HotelMsg*, C uses r to send a message to H for asking for a quote. Then, the hotel will answer with an offer on port *offer*. Whether the customer accepts or not a particular quote is modeled by the multiple definitions for the pattern *offer*($\$$) in C . If the offer is not adequate then C can abort the negotiation, which will activate the compensation $Q \mid Q'$ (analogously for H and A). C can accept the offer by sending a confirmation message on port c . In this case, C also generates a message to the local port *hotelOK*($hConf$). This message will be managed by the local merge rule defined by C . The contract will be blocked until a running negotiation for buying flight tickets generates a message on *flightOK*. At this time, the local merge definition can be fired and both contracts merged. Eventually the negotiation will commit by releasing the messages on *HotelFound* and *FlightFound*. Moreover, messages *BookedRoom*(v) and *SoldFlight*(v) generated by H and A to change their local states are released only at this time, when all participants have committed.

Andorra Kernel Language. As a second example, we sketch how merge definitions and nesting can be used to model some features of AKL [13], a concurrent logic programming language. We consider guarded rules $A \leftarrow G \mid B$, where the *head* A is an atom, the *guard* G and the *body* B are (possibly empty) conjunction of atoms, and \mid is the commit operator. An AKL program \mathcal{P} is a list of guarded rules. An execution of \mathcal{P} is initiated by providing a *goal* \mathcal{G} , which is a conjunction of atoms. A **cJoin** process that simulates \mathcal{P} queried with \mathcal{G} is:

$$\mathbf{def} \ D \wedge \llbracket \mathbf{defs}(\mathcal{P}) \rrbracket \wedge \llbracket \mathbf{undef}(\mathcal{P}) \rrbracket \ \mathbf{in} \ \llbracket \mathbf{and}(\mathcal{G}) \rrbracket_{trueG, falseG} \mid \mathbf{unif}\langle final, \mathbf{tt} \rangle$$

The definitions in D are needed to promote constraints computed locally. The rules in \mathcal{P} are translated separately by grouping all rules defining the same atom A . Such partition is denoted by $\mathbf{defs}(\mathcal{P})$, while $\mathbf{undef}(\mathcal{P})$ denotes all atoms in \mathcal{P} without defining rules, i.e., atoms whose proofs will always fail. Constraints θ are encoded conveniently as **cJoin** processes (\mathbf{tt} is the empty constraint). In general, the term $\mathbf{unif}\langle final, \theta \rangle$ stores the computed constraints of a running proof. At the end, a message on either port *trueG* or *falseG* will inform the environment about the outcome of the computation. The encoding of clauses and goals is in Figure 6.

An atom A is encoded as a merge rule that substitutes a message $A\langle t, f \rangle$ by the proof of its defining rules (rule **DEF**). An undefined atom is encoded as a rule that always fails (**UNDEF**). A conjunction (**AND**) corresponds to a process that activates a new negotiation containing the atoms $A_i\langle t_i, f_i \rangle$ to be proved and the initial local constraints $\mathbf{unif}\langle w, \mathbf{tt} \rangle$. Every sub-proof initiated by $A_i\langle t_i, f_i \rangle$ will notify its termination by using ports t_i (success) and f_i (failure). If all sub-proofs end successfully, the second definition in the negotiation can be fired producing $t\langle \rangle$ (i.e., the signal of the successful proof of the conjunction) and

$$\begin{array}{l}
(\text{DEF}) \quad \llbracket A \leftarrow B_1 | C_1, \dots, A \leftarrow B_n | C_n \rrbracket = A\langle t, f \rangle \blacktriangleright \llbracket \text{choice}(B_1 | C_1, \dots, B_n | C_n) \rrbracket_{t,f} \\
(\text{UNDEF}) \quad \llbracket A \rrbracket = A\langle t, f \rangle \blacktriangleright f\langle \rangle \\
(\text{AND}) \quad \llbracket \text{and}(A_1, \dots, A_n) \rrbracket_{t,f} = \text{def } \text{and}\langle \rangle \triangleright [\text{def } w\langle \rangle \triangleright 0 \\
\quad \quad \quad \wedge \prod_{i=1}^n t_i\langle \rangle \triangleright \text{done}\langle w \rangle | t\langle \rangle \\
\quad \quad \quad \wedge \prod_{i=1}^n f_i\langle \rangle \triangleright \text{and}\langle \rangle | \text{failed}\langle \rangle \\
\quad \quad \quad \dots \text{Rules for handling abortion...} \\
\quad \quad \quad \text{in } \prod_{i=1}^n A_i\langle t_i, f_i \rangle | \text{unif}\langle w, \text{tt} \rangle : f\langle \rangle] \\
\text{in } \text{and}\langle \rangle \\
(\text{CHOICE}) \llbracket \text{choice}(A_1 | B_1, \dots, A_n | B_n) \rrbracket_{t,f} = \text{def } \wedge_{i=1}^n i\langle \rangle \triangleright \llbracket \text{and}(B_i) \rrbracket_{t,f} \\
\quad \quad \quad \text{in } [\text{def } \prod_{i=1}^n f_i\langle \rangle | \text{proving}\langle \rangle \triangleright \text{abort} \\
\quad \quad \quad \wedge_{i=1}^n t_i\langle \rangle | \text{proving}\langle \rangle \triangleright i\langle \rangle | \text{chosen}\langle \rangle \\
\quad \quad \quad \dots \text{Rules for handling abortion...} \\
\quad \quad \quad \text{in } \prod_{i=1}^n \llbracket \text{and}(A_i) \rrbracket_{t_i, f_i} \text{proving}\langle \rangle : f\langle \rangle]
\end{array}$$

Figure 6. Encoding of AKL committed rules.

$\text{done}\langle w \rangle$ that activates the promotion of computed constraints managed by D (omitted for space limitation). Note that $t\langle \rangle$ will be released outside only at commit, after constraints have been promoted. Instead, if a sub-atom fails, all running sub-contracts are aborted and the contract commits by releasing the activation of a new proof.

Rule **CHOICE** opens a negotiation for proving one of the guards of a multiple choice goal. When a guard is successful the negotiation can commit by releasing the message $i\langle \rangle$, which activates the body B_i of the chosen goal.

If there is an AKL refutation for the goal \mathcal{G} with computed constraints $\text{and}(\text{tt}, \theta)$ then the messages $\text{final}\langle \theta \rangle$ and $\text{true}\mathcal{G}\langle \rangle$ can be released. On the other hand, $\text{false}\mathcal{G}\langle \rangle$ is generated only if \mathcal{G} cannot be proved.

5. Encoding Zero-Safe nets

Zero-safe nets (*zs nets*) [8] extends Place/Transition Petri nets (*PT nets*) with a mechanism for expressing concurrent transactions. Recently, they have been used in [7] to encode short-running transactions of Biztalk, a commercial workflow management system [16]. *ZS* nets additionally provides a “dynamic” specification of transactions boundaries supporting multiway transactions, which retain several entry and exit points, and admit a number of participants which is statically unknown. However, *ZS* nets are not suitable to express some interesting aspects, such as mobility, programmable compensations and nesting.

In this section we show that *ZS* nets can be straightforwardly encoded in **cJoin**. A distributed implementation of *ZS* nets in **Join** has been presented in [7], but there the encoding is complicated by the need of attaching a local transaction manager to each transition.

We recall that, in Petri nets, *places* are repositories of *tokens* and *transitions* fetch and produce tokens. Net configurations, called *markings*, are multisets of tokens. The places of *ZS* nets are partitioned into ordinary and transactional

$$\begin{array}{c}
\text{(FIRING)} \\
\frac{S + Z \mid \rangle \quad S' + Z' \in T}{(S + S'', Z + Z'') \rightarrow_T (S' + S'', Z' + Z'')} \\
\text{(CONCATENATION)} \\
\frac{(S_1, Z) \rightarrow_T (S'_1, Z'') \quad (S_2, Z'') \rightarrow_T (S'_2, Z')}{(S_1 + S_2, Z) \rightarrow_T (S'_1 + S'_2, Z')} \\
\text{(STEP)} \\
\frac{(S_1, Z_1) \rightarrow_T (S'_1, Z'_1) \quad (S_2, Z_2) \rightarrow_T (S'_2, Z'_2)}{(S_1 + S_2, Z_1 + Z_2) \rightarrow_T (S'_1 + S'_2, Z'_1 + Z'_2)} \\
\text{(CLOSE)} \\
\frac{(S, \emptyset) \rightarrow_T (S', \emptyset)}{(S, \emptyset) \Rightarrow_T (S', \emptyset)}
\end{array}$$

Figure 7. Operational semantics of zs nets (+ denotes multiset union).

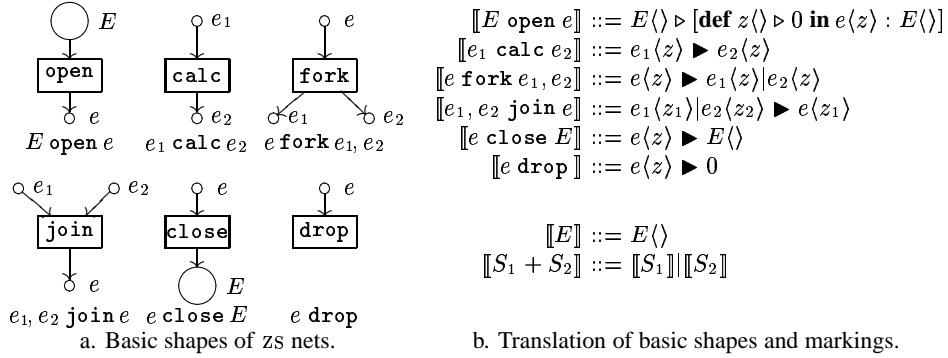


Figure 8. Encoding of zs nets in cJoin.

ones, called *stable* and *zero*, respectively. Correspondingly, markings U can be seen as pairs (S, Z) with $U = S + Z$, where S and Z are the multisets of stable and zero resources, respectively. Transitions are written $U \mid \rangle U'$. A *transaction* goes from a multiset of stable places (*stable marking*) to another stable marking. The key point is that stable tokens produced during a transaction are made available only at commit time, when no zero tokens are left. We write (T, S) for a zs net with set of transitions T and initial stable marking S .

The operational semantics of zs nets is defined by the two relations \Rightarrow_T and \rightarrow_T in Figure 7. Rules FIRING and STEP are the ordinary ones for Petri nets, for the execution of one/many transition(s). However, sequences of steps differ from the ordinary transitive closure of \rightarrow_T : The rule CONCATENATION composes zero tokens in series but stable tokens in parallel, hence stable tokens produced by the first step cannot be consumed by the second step. CLOSE selects the moves $(S, \emptyset) \Rightarrow_T (S', \emptyset)$, which defines the transactions of the net.

As done in [7], and without loss of generality, we restrict to zs nets whose transitions have the basic shapes in Figure 8.a, for E any stable place and e, e_1, e_2 any zero places. The translation in Figure 8.b associates a cJoin definition to each basic shape. Places are seen as ports and tokens as messages. Tokens in stable places carry no value, while tokens in zero places carry the identifier of the transaction they belong to. The cJoin process associated to the zs net (T, S) is $\mathbf{def} \llbracket T \rrbracket \mathbf{in} \llbracket S \rrbracket$, where $\llbracket T \rrbracket = \bigwedge_{t \in T} \llbracket t \rrbracket$.

A transaction can be opened by firing a transition of the form $E \text{ open } e$. In cJoin , this means opening a new negotiation whose internal state contains the definition of a fresh name z , the message $e\langle z \rangle$, and whose compensation, by default, gives back the stable resources $E\langle \rangle$. The dummy definition $z\langle \rangle \triangleright 0$ is the cJoin way of declaring a fresh identifier z for the transaction. When two transactions are merged by applying e.g. $e_1\langle z_1 \rangle | e_2\langle z_2 \rangle \blacktriangleright e\langle z_1 \rangle$, then z_1 and z_2 become equivalent identifiers for the same larger negotiation. When computing inside a negotiation, each zero token carries one of the possibly many equivalent identifiers for that negotiation (e.g., z_1). If stable messages $E\langle \rangle$ are released inside the negotiation, e.g., by firing $e \text{ close } E$, then they are frozen until commit, because the only rules that can fetch them are outside the negotiation boundaries, in the top chemical soup. The commit can happen if and only if the negotiation reaches a local state containing only stable messages (and dummy definitions). Then, the reaction COMMIT can close the negotiations and release all stable tokens to the environment. The following result assures the correctness and completeness of $\llbracket - \rrbracket$.

THEOREM 2 $(S, \emptyset) \Rightarrow_T^* (S', \emptyset)$ iff $\mathbf{def} \llbracket T \rrbracket \mathbf{in} \llbracket S \rrbracket \rightarrow^* \mathbf{def} \llbracket T \rrbracket \mathbf{in} \llbracket S' \rrbracket$.

The main behavioral difference between the cJoin encoding in Figure 8 and the Join encoding in [7] relies on the treatment of failures, as here no abort can be generated (and consequently compensations cannot be activated). A possible solution would be to add a *timeout* component each time a new negotiation is open, which is able to produce the abort via the rule $A = \text{timeout}\langle \rangle \blacktriangleright \text{abort}$. In this case, we should let $\llbracket E \text{ open } e \rrbracket = E\langle \rangle \triangleright [\mathbf{def} z\langle \rangle \triangleright 0 \mathbf{in} e\langle z \rangle | \text{timeout}\langle \rangle : E\langle \rangle]$ and encode the net $N = (T, S)$ as $\mathbf{def} A \wedge \llbracket T \rrbracket \mathbf{in} \llbracket S \rrbracket$.

6. Serializability and Big-Step Semantics

The semantics of cJoin given in Figure 4 allows the cooperation among several negotiations. Nevertheless, we would like to reason about a process by analyzing interacting negotiations independently from the rest of the system. A concurrent execution $T_1 \parallel \dots \parallel T_n$ of several transactions is said *serializable* if there exists a sequence $T_{i_1}; T_{i_2}; \dots; T_{i_n}$ that executes all transactions one at a time (without interleaving their steps) and produces the same result [2]. Serializability is important because it allows to reason about the behavior of a system by considering one transaction at a time. In this section we introduce a syntactical restriction on processes, called *shallowness*, and show that it guarantees serializability.

The idea is to describe multi-party negotiations as abstract transitions that fetch the messages needed to initiate all sub-negotiations separately and produce the processes released at commit or abort. Consequently, serializable negotiations can postpone the activation of each sub-negotiation until all other cooperating sub-negotiations needed to commit can be activated.

$$\begin{array}{c}
\text{(PAR)} \\
\frac{\mathcal{D} \vdash \mathcal{P} \rightarrow \mathcal{D} \vdash \mathcal{P}' \quad \mathcal{D} \vdash \mathcal{Q} \rightarrow \mathcal{D} \vdash \mathcal{Q}'}{\mathcal{D} \vdash \mathcal{P} \mid \mathcal{Q} \rightarrow \mathcal{D} \vdash \mathcal{P}' \mid \mathcal{Q}'} \\
\text{(GLOBAL FIRING)} \\
\mathcal{D} \wedge J \triangleright \mathcal{P} \vdash J\sigma \rightarrow \mathcal{D} \wedge J \triangleright \mathcal{P} \vdash \mathcal{P}\sigma \\
\text{(MERGE)} \\
\mathcal{D} \wedge \Pi_i J_i \blacktriangleright \mathcal{S} \vdash \Pi_i [\mathcal{D}_i \vdash J_i \sigma \mid \mathcal{S}_i : \mathcal{Q}_i] \rightarrow \mathcal{D} \wedge \Pi_i J_i \blacktriangleright \mathcal{S} \vdash [\bigwedge_i \mathcal{D}_i \vdash (\Pi_i \mathcal{S}_i) \mid \mathcal{S}\sigma : \Pi_i \mathcal{Q}_i] \\
\text{(LOCAL COMMIT)} \quad \text{(ABORT)} \quad \text{(IDLE)} \\
\mathcal{D} \vdash [M \mid \mathcal{D}' \vdash 0 : \mathcal{S}] \rightarrow \mathcal{D} \vdash M \quad \mathcal{D} \vdash [\text{abort} \mid \mathcal{P} : \mathcal{S}] \rightarrow \mathcal{D} \vdash \mathcal{S} \quad \mathcal{D} \vdash \mathcal{P} \rightarrow \mathcal{D} \vdash \mathcal{P}
\end{array}$$

Figure 9. Big-step semantics of cJoin.

DEFINITION 3 (SHALLOWNESS) *The nesting $nest(P)$ of P is defined by:*

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

P is shallow if any basic definition D in P satisfies one of the two conditions:

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

We refer definitions in shallow processes as *shallow definitions*. Moreover, we call a process P *stable* iff P is shallow and $nest(P) = 0$. The shallow property imposes a discipline for activating negotiations. In particular, condition 1 assures that the firing of an ordinary rule increases the height of the nesting by at most one level (i.e., a definition produces either a stable process or one negotiation without nested sub-contracts). Condition 2 forbids the creation of sub-negotiations while merging. The absence of condition 2 would prevent the possibility of postponing the activation of some negotiations until all cooperating negotiations can be activated. Shallowness forbids rules such as $J_1 \triangleright P \mid [P_1 : Q]$ and $J_2 \triangleright [[P_1 : Q_1] : Q]$, which however can be encoded as shallow definitions by using new local ports, e.g. as $D_1 \equiv J_1 \triangleright x\langle \rangle \mid P \wedge x\langle \rangle \triangleright [P_1 : Q]$ and $D_2 \equiv J_2 \triangleright [\mathbf{def} x\langle \rangle \triangleright [P_1 : Q_1] \mathbf{in} x\langle \rangle : Q]$, respectively. Moreover $\mathbf{def} D_2 \mathbf{in} J_2\sigma$ reduces in two steps to $\mathbf{def} D_2 \mathbf{in} [\mathbf{def} x\langle \rangle \triangleright [P_1\sigma : Q_1\sigma] \mathbf{in} [P_1\sigma : Q_1\sigma] : Q\sigma]$, which has nested negotiations.

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 $\mathbf{def} \mathcal{D} \mathbf{in} \mathcal{P}$ 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

$$\begin{aligned}
\mathcal{D} \vdash \mathcal{P} \mid \mathbf{def} \mathcal{D}' \mathbf{in} \mathcal{Q} &\equiv \mathcal{D} \wedge \mathcal{D}' \sigma_{dn} \vdash \mathcal{P} \mid \mathcal{Q} \sigma_{dn} \\
range(\sigma_{dn}) \cap (fn(\mathcal{D}) \cup fn(\mathcal{P}) \cup fn(\mathbf{def} \mathcal{D}' \mathbf{in} \mathcal{Q})) &= \emptyset
\end{aligned}$$

We characterize serializability by the big-step reduction relation between shallow processes presented in Figure 9.

Steps can be composed in parallel (PAR) and sequentially (SEQ), even with idle transitions (IDLE). Rule GLOBAL FIRING corresponds to the firing of an ordinary definition in a top-level process. Instead LOCAL FIRING states possible internal transitions of a running contract. LOCAL FIRING represents suitable sub-negotiations as ordinary transitions at an abstract level. In fact, the computations occurring at a lower level in the nesting hierarchy (premise of LOCAL FIRING) that are relevant to its containing negotiation are those relating stable processes, i.e., \mathcal{S} and \mathcal{S}' . A negotiation has available, in addition to its own definitions, the merge definitions introduced by its parent. In fact, a merge definition applied on a single contract behaves as an ordinary rule but defined in a global scope. The operator $\widetilde{}$ transforms merge definitions in ordinary ones: $\widetilde{J \triangleright \mathcal{P}} \equiv J \triangleright \mathcal{P}$ and $\widetilde{\mathcal{B} \wedge \mathcal{B}'} \equiv \widetilde{\mathcal{B}} \wedge \widetilde{\mathcal{B}'}$. If the rule STR-MOVE is also considered for cJoin, then the premise of LOCAL FIRING must be $\widetilde{\mathcal{B}} \wedge \mathcal{B} \vdash \mathcal{S} \rightarrow \widetilde{\mathcal{B}} \wedge \mathcal{B} \vdash \mathcal{S}'$.

Rules LOCAL COMMIT and ABORT handle the termination of a negotiation, whereas MERGE describes the interaction among sibling negotiations. This time, negotiations can be joined only if they do not contain running contracts.

The big-step relation enforces serializability. In fact, the completed negotiations at a particular level become ordinary transitions at the upper level and all interacting transactions can be analyzed independently from the rest of the system. The following result states the correspondence between both semantics for shallow processes, proving that shallow processes are serializable.

THEOREM 4 *Let $\mathcal{S}, \mathcal{S}'$ be stable processes. Then $\mathcal{S} \rightarrow^* \mathcal{S}'$ iff $\vdash \mathcal{S}' \rightarrow \vdash \mathcal{S}'$.*

7. Concluding remarks

We have proposed cJoin as a linguistic extension of Join with natural primitives for modeling nested negotiations. The expressiveness of cJoin has been demonstrated by means of an informal discussion about the satisfaction of the general requirements enumerated in the Introduction and three sample applications (trip booking and the encodings of AKL and ZS nets). Additionally, we have defined a syntactical restriction of processes that assures serializability.

Unlike *workflow systems*, cJoin does not fix a set of constructors to describe dependencies in the style of ConTracts [15]. The actual dependencies of a negotiation are known at execution time and are consequence of its interaction with other contracts. This feature distinguishes cJoin from [4], where processes running as transactions can interact freely with the environment. On the other hand, cJoin is aimed at providing a way to model multi-party transactions by describing their interacting agents and not their global structure, such as [6, 9]. Nevertheless, our language does not provide default mechanisms for undoing pre-committed activities of aborted transactions, differently from [9]. We leave as a future work the comparison with other calculi that models ACID transactions [17], long-running negotiations [4, 11], and exception handling [14].

Negotiations in **cJoin** have the flavour of multiway transactions in **ZS** nets, where participants are not statically fixed. We plan to reuse or to extend the **D2PC** proposed in [7] to have a full encoding of **cJoin** in **Join** itself. This would allow us to extend implementations of **Join**, such as **Jocaml** or **Polyphonic C[#]**, by providing primitives for handling distributed negotiations. As a preliminary result, the encoding for the subcalculus of *flat* processes is in [5].

Acknowledgments. We thank Nick Benton, Luca Cardelli, Cédric Fournet and Cosimo Laneve with whom we discussed preliminary versions of **cJoin**.

References

- [1] N. Benton, L. Cardelli, and C. Fournet. Modern concurrency abstractions for **C[#]**. *Proc. of ECOOP 2002, LNCS 2374*, pp. 415–440. Springer Verlag, 2002.
- [2] P.A. Bernstein, V. Hadzilacos, and N. Goodman. *Concurrency, Control and Recovery in Database Systems*. Addison-Wesley Longman, 1987.
- [3] G. Berry and G. Boudol. The chemical abstract machine. *TCS*, 96(1):217–248, 1992.
- [4] L. Bocchi, C. Laneve, and G. Zavattaro. A calculus for long-running transactions. *Proc. of FMOODS'03, LNCS 2884*, pp. 194–208. Springer Verlag, 2003.
- [5] R. Bruni, H. Melgratti, and U. Montanari. Flat Committed Join in Join. *Proc. of COMETA 2003, ENTCS*. To appear.
- [6] BPEL Specification, May 2003. <http://www-106.ibm.com/developerworks/webservices/library/ws-bpel/>.
- [7] R. Bruni, C. Laneve, and U. Montanari. Orchestrating transactions in join calculus. *Proc. of CONCUR 2002, LNCS 2421*, pp. 321–336. Springer Verlag, 2002.
- [8] R. Bruni and U. Montanari. Zero-safe nets: Comparing the collective and individual token approaches. *Inform. and Comput.*, 156(1-2):46–89, 2000.
- [9] M. Butler, M. Chessell, C. Ferreira, C. Griffin, P. Henderson, and D. Vines. Extending the concept of transaction compensation. *IBM Systems Journal*, 41(4):743–758, 2002.
- [10] S. Conchon and F. Le Fessant. **Jocaml**: Mobile agents for Objective-Caml. *Proc. of ASA'99/MA'99*, 1999.
- [11] D. Duggan. Abstractions for Fault-Tolerant Global Computing. *TCS*. To appear.
- [12] C. Fournet and G. Gonthier. The reflexive chemical abstract machine and the Join calculus. *Proc. of POPL'96*, pp. 372–385. ACM Press, 1996.
- [13] S. Haridi, S. Janson, and C. Palamidessi. Structural operational semantics of **AKL**. *Journal of Future Generation Computer Systems*, 8:409–421, 1992.
- [14] M. Mazzara and R. Lucchi. A framework for generic error handling in business processes. *Proc. of WS-FM'04, ENTCS*. To appear.
- [15] A. Reuter and H. Wächter. The contract model. *Transaction Models for Advanced Applications*. Morgan Kaufmann, 1992.
- [16] U. Roxburgh. Biztalk orchestration: transactions, exceptions, and debugging, 2001. <http://msdn.microsoft.com/library/en-us/dnbiz/html/btsorch.asp>.
- [17] J. Vitek, S. Jagannathan, A. Welc, and A. L. Hosking. A Semantic Framework for Designer Transactions. *Proc. of ESOP'04, LNCS 2986*, pp. 249–263. Springer Verlag, 2004.