

Choreography Synthesis as Contract Agreement

Julien Lange

University of Leicester, UK

`jlange@le.ac.uk`

Alceste Scalas

University of Cagliari, Italy

`alceste.scalas@unica.it`

We propose a formal model for distributed systems, where each participant advertises its requirements and obligations as behavioural *contracts*, and where multiparty sessions are started when a set of contracts allows to synthesise a *choreography*. Our framework is based on the CO₂ calculus for contract-oriented computing, and borrows concepts and results from the session type literature.

It supports sessions where the number of participants is not determined beforehand, and keeps CO₂'s ability to rule out participants that are *culpable* if contracts are not fulfilled at runtime. We show that we have *progress* and *session fidelity* in CO₂, as a result of the *honesty* of participants — i.e., their ability to always adhere to their contracts.

1 Introduction

Distributed applications are nowadays omnipresent but even for seemingly simple cases, there is still a pressing need to make sure they do work as their designers intended. Indeed, such systems are difficult to design, verify, implement, deploy, and maintain. Besides the intrinsic issues due to the underlying execution model (concurrency, physical distribution, etc.), applications have to be designed within a strange paradox: they are made of components that, on the one hand, collaborate with each other and, on the other hand, may compete for resources, or for achieving conflicting goals. This paradox is especially relevant in inter-organisational service-oriented scenarios, where services may be deployed by different entities: even under common policies, the implementations may reflect diverging and changing requirements, up to the point of departing from the agreed specifications. This issue is reflected in standards such as [15], which includes runtime monitoring and logging to check that interactions in SOAs actually adhere to agreed policies and service descriptions.

Along the lines of [5, 4], we propose a formal model for distributed systems where *contracts* drive interactions: components advertise behavioural contracts; such contracts are used at runtime to establish multiparty *agreements*, and such agreements steer the behaviour of components. Therefore, contracts are not just a specification or a design mechanism anymore, rather they become a pivotal element of the execution model.

In this work, we combine two approaches: *session types* [11] and *contract-oriented computing* [5]. From the former, we adopt concepts, syntax and semantics — and in particular, the interplay between local behaviours and choreographies (i.e., between local types and global types) as a method for specifying and analysing the interactions of participants in a distributed system. However, in our framework we do not assume that a participant will necessarily always adhere to its specification, nor that a global description is available beforehand to validate the system. From the contract-oriented computing approach, we adopt CO₂ [4], a generic contract-oriented calculus where participants advertise their requirements and obligations through contracts, and interact with each other once *compliant* contracts have been found. Here, we tailor CO₂ to a multiparty model where contracts have the syntax of local types. We say that contracts c_1, \dots, c_n are compliant when, roughly, they can be used to synthesise a choreography — i.e., a global type whose projections are c_1, \dots, c_n themselves [14]. Once a set of compliant contracts has

been found, a CO₂ session may be established, wherein the participants who advertised the contracts can interact. However, in line with what may happen in real life scenarios, the runtime behaviour of these participants may then depart from the contracts: the calculus allows to model these situations, and reason about them.

1.1 Contributions

Our framework models multiparty contractual agreements as “tangible” objects, i.e., choreographies. This allows us to rely on results and properties from the session type literature — in particular, the well-formedness of a choreography ensures that contractual agreements enjoy knowledge of choice, error/deadlock freedom, and progress. Furthermore, it allows us to easily check that some meta-level properties are satisfied at runtime, e.g., on the number of involved participants, whether or not the session may terminate, etc.

Our adaptation of CO₂ to a multiparty, choreography-based contract model preserves the properties of the original calculus. In particular, if a system gets stuck, it is possible to identify which participants violated their contracts.

We also discuss how the properties of a well-formed choreography are reflected in a context where participants can misbehave. We introduce global progress and session fidelity in CO₂, again inspired by analogous concepts in theories based on session types. We show that they hold in systems where all participants are *honest* (i.e., they always respect their contracts in any context) — even when a participant takes part in multiple sessions.

Synopsis. The rest of the paper is structured as follows. In the rest of this section, we introduce an example that we use to motivate and illustrate our framework. In Section 2, we introduce a multiparty contract model based on choreography synthesis. In Section 3, we present our version of CO₂ and highlight its main features. In Section 4, we define the notion of *honesty* and its practical importance in our contract-oriented scenario. In Section 5, we present our results, which link the notion of honesty to the progress and safety of a CO₂ system (due to lack of space, the proofs are relegated to an online appendix [13]). Finally, we discuss related work and conclude in Section 6.

1.2 A motivating example

In this section, we introduce a running example to illustrate our framework. We use A, B, \dots for participant names, and a, a', b, \dots for participant variables, and use the colour **blue** to highlight contracts.

Consider the following distributed scenario: an online store A allows two buyers b_1 and b_2 to make a joint purchase through a simplified protocol: after they both request the same item, a quote is sent to b_1 , who is then expected to either place an order (`order`) or end the session (`bye`); the store also promises to notify b_2 about whether the order was placed (`ok`) or cancelled (`bye`). A 's behaviour is described by the following contract:

$$c_A = \mathbf{b_1?req; b_2?req; b_1!quote; (b_1?order; b_2!ok + b_1?bye; b_2!bye)}$$

What kind of contracts would be compliant with c_A ? One answer consists in the following contracts, advertised by buyers B_1 and B_2 .

$$\begin{aligned} c_{B_1} &= \mathbf{a!req; a?quote; (b'_2!ok; a!order \oplus b'_2!bye; a!bye)} \\ c_{B_2} &= \mathbf{a'!req; (b'_1?ok; a'?ok + b'_1?bye; a'?bye)} \end{aligned}$$

Here, B_1 promises to send the request to the store (a), wait for the quote, and then notify the other buyer (b'_2) before accepting or rejecting the store offer; symmetrically, B_2 's contract sends the request to the store (a'), and then expects to receive the same notification (either ok or bye) from both the other buyer (b'_1) and the store itself. Each contract represents the local viewpoint of the participant who advertises it: c_A represents the local viewpoint of the store, and thus it does not (and indeed, it cannot) capture the communications between B_1 and B_2 .

An agreement among c_A , c_{B_1} and c_{B_2} may be found by replacing the participant variables in each contract with actual names, e.g., with substitutions $\{A/a, a'\}$, $\{B_1/b_1, b'_1\}$ and $\{B_2/b_2, b'_2\}$. Such an agreement is based on the existence of the following choreography (i.e., global type), which can be synthesised similarly to what is done in [14]:

$$\begin{aligned} \mathcal{G}_{AB_1B_2} = & B_1 \rightarrow A : req ; B_2 \rightarrow A : req ; A \rightarrow B_1 : quote ; \\ & (B_1 \rightarrow B_2 : ok ; B_1 \rightarrow A : order ; A \rightarrow B_2 : ok \quad + \quad B_1 \rightarrow B_2 : bye ; B_1 \rightarrow A : bye ; A \rightarrow B_2 : bye) \end{aligned}$$

The ability to synthesise $\mathcal{G}_{AB_1B_2}$ guarantees that the global type is well-formed and projectable back to the initial contracts c_A , c_{B_1} and c_{B_2} (with the substitutions above); this, in turn, guarantees progress and safety [14] of the contractual agreement.

However, in a realistic scenario, the existence of a contractual agreement among participants does not guarantee that progress and safety will also hold at runtime: in fact, a participant may advertise a contract promising some behaviour, and then fail to respect it — either maliciously or accidentally. Such failure may then cascade on other participants, e.g., if they remain stuck waiting for a promised message that is never sent.

This sort of situations can be modelled using the CO_2 calculus. A CO_2 system for the store-and-two-customers example may be implemented as follows:

$$S_1 = (x, y, z) (A[\text{tell}_A \downarrow_x c_A \cdot \text{fuse} \cdot P_A] \mid B_1[\text{tell}_A \downarrow_y c_{B_1} \cdot P_{B_1}] \mid B_2[\text{tell}_A \downarrow_z c_{B_2} \cdot P_{B_2}])$$

Here, participant A advertises its contract c_A to itself via the primitive $\text{tell}_A \downarrow_x c_A$, where x is used as a session handle for interacting with other participants. B_1 and B_2 advertise their respective contracts to A with a similar invocation.

In this example, A also plays the role of *contract broker*: once all contracts have been advertised, the fuse prefix can establish a new session, based on the fact that the global agreement $\mathcal{G}_{AB_1B_2}$ can be synthesised from c_A , c_{B_1} and c_{B_2} . This new session is shared among participants A , B_1 and B_2 .

At this point, the execution of the system (i.e., the continuation of processes P_A , P_{B_1} and P_{B_2}) is not required to respect the contracts. In fact, we will see that when the contracts are violated, the calculus allows for *culpable* participants to be always ruled out. Furthermore, we will discuss *honesty*, i.e., the guarantee that a participant will always fulfil its advertised contracts — even in contexts where other participants fail to fulfil theirs. When such a guarantee holds, the contractual progress and safety are also reflected in the runtime behaviour of the CO_2 system.

Other possible agreements. Our contract model allows for other scenarios. For instance, a participant B_{12} may impersonate both customers, and promise to always accept the store offer, by advertising the following contract:

$$c_{B_{12}} = a''!req; a''!req; a''?quote; a''!order; a''?ok$$

where the request to the store (a'') is sent twice (i.e., once for each impersonated customer). In this case, if we combine c_A and $c_{B_{12}}$ with substitutions $\{A/a''\}$, $\{B_{12}/b_1, b_2\}$, we can find an agreement by synthesising

the following global type:

$$\mathcal{G}_{AB_{12}} = B_{12} \rightarrow A : \text{req}; B_{12} \rightarrow A : \text{req}; A \rightarrow B_{12} : \text{quote}; B_{12} \rightarrow A : \text{order}; A \rightarrow B_{12} : \text{ok}$$

Similarly to the previous case, this scenario may be modelled with the following CO₂ system:

$$S_2 = (x, w) (A[\text{tell}_A \downarrow_x c_A \cdot \text{fuse} \cdot P_A] \mid B_{12}[\text{tell}_A \downarrow_w c_{B_{12}} \cdot P_{B_{12}}])$$

where the fuse prefix can now create a session involving A and B₁₂.

The participants in the CO₂ systems S_1 and S_2 may also be combined, so to obtain:

$$S_{12} = (x, y, z, w) (A[\text{tell}_A \downarrow_x c_A \cdot \text{fuse} \cdot P_A] \mid B_1[\text{tell}_A \downarrow_y c_{B_1} \cdot P_{B_1}] \mid B_2[\text{tell}_A \downarrow_z c_{B_2} \cdot P_{B_2}] \\ \mid B_{12}[\text{tell}_A \downarrow_w c_{B_{12}} \cdot P_{B_{12}}])$$

In this case, after all contracts have been advertised to A, either a session corresponding to $\mathcal{G}_{AB_1B_2}$, or to $\mathcal{G}_{AB_{12}}$ may take place, thus involving a different number of participants depending on which contracts are fused. In such cases, it makes sense to consider whether one of the agreements should take precedence over the other, and which criteria should drive this choice.

2 A Choreography-Based Contract Model

We introduce a contract model based on concepts and results from the session types literature. Individual contracts are expressed using the syntax of local session types; while contractual compliance is based on global types synthesis: a set of contracts is compliant if it is possible to synthesise a choreography from it, as described in [14]. For simplicity, we adopt syntax and semantics in the style of [8, 10]: we use participant names (instead of channels) for message exchange, i.e., we consider systems with just one channel between each pair of participants.

Syntax & Semantics. Let \mathbb{P} and \mathcal{P} be disjoint sets of, respectively, *participant names* (ranged over by A, B, \dots) and *participant variables* (ranged over by a, b, \dots). Let $\underline{a}, \underline{b}$ range over $\mathbb{P} \cup \mathcal{P}$. The syntax of contracts below is parametrised wrt *sorts* (ranged over by e) which abstract data types (either simple or complex). We use the colour **blue** for single contracts and **green** for *systems* of contracts.

$$\begin{array}{lcl} T, T' & ::= & T \mid T' \quad \mid \quad A\langle c \rangle \quad \mid \quad (AB) : \rho \quad \mid \quad \mathbf{0} \\ c, c' & ::= & \bigoplus_{i \in I} \underline{a}_i ! e_i ; c_i \quad \mid \quad \sum_{i \in I} \underline{a} ? e_i ; c_i \quad \mid \quad \mu \mathbf{x}. c \quad \mid \quad \mathbf{x} \end{array}$$

A contract c may be either: (i) an internal choice \bigoplus , with the intuitive semantics that after sending the message e_i to participant \underline{a}_i , behaviour c_i take places; (ii) an external choice \sum , saying that if a message of sort e_i is received from \underline{a} , then behaviour c_i takes place; or (iii) a recursive behaviour. We assume that $\forall i \neq j \in I: (\underline{a}_i, e_i) \neq (\underline{a}_j, e_j)$ in internal choices, $\forall i \neq j \in I: e_i \neq e_j$ in external choices, and that \bigoplus and \sum are associative and commutative. We write $\text{fv}(c)$ for the free participant variables in c .

A system of contracts T may be either: (i) a parallel composition of systems $T \mid T'$; or (ii) a *named* contract $A\langle c \rangle$, saying that participant A promises to behave according to c ; (iii) a *queue* $(AB) : \rho$ of messages from A to B. In a system T , we assume that there is at most one queue per pair of participants, (i.e., one channel per direction), and that participant names are pairwise distinct.

We consider systems of contracts as processes whose semantics is given by the following main reduction rules (see the online appendix [13] for the omitted ones):

$$\begin{array}{l} A\langle B!e; c_0 \oplus c_1 \rangle \mid (AB):\rho \mid T \xrightarrow{A \rightarrow B:e} A\langle c_0 \rangle \mid (AB):\rho \cdot e \mid T \\ A\langle B?e; c_0 + c_1 \rangle \mid (BA):e \cdot \rho \mid T \xrightarrow{A \leftarrow B:e} A\langle c_0 \rangle \mid (BA):\rho \mid T \end{array}$$

The first rule says that, after an internal choice, participant A puts a message e on its queue for participant B. The second rule says that A's external choice can receive a message of the right sort from an input queue BA. We write $T \xrightarrow{A \leftarrow B:e} T'$ when either $T \xrightarrow{A \rightarrow B:e} T'$ or $T \xrightarrow{A \leftarrow B:e} T'$, and $Q(T)$ for the parallel composition of the empty queues connecting all pairs of participants in T .

Example 2.1. *From the example in Section 1.2, consider the instantiated contracts of the store A and its customer B_{12} . We illustrate the initial system, and how it progresses:*

$$\begin{array}{l} T_{AB_{12}Q} = A\langle c_A \{A/a''\} \{B_{12}/b_1, b_2\} \rangle \mid B_{12}\langle c_{B_{12}} \{A/a''\} \{B_{12}/b_1, b_2\} \rangle \mid (AB_{12}):[] \mid (B_{12}A):[] \\ = A\langle B_{12}?req; B_{12}?req; \dots \rangle \mid B_{12}\langle A!req; A!req; \dots \rangle \mid (AB_{12}):[] \mid (B_{12}A):[] \\ \xrightarrow{B_{12} \rightarrow A:req} A\langle B_{12}?req; B_{12}?req; \dots \rangle \mid B_{12}\langle A!req; \dots \rangle \mid (AB_{12}):[] \mid (B_{12}A):req \\ \xrightarrow{A \leftarrow B_{12}:req} A\langle B_{12}?req; \dots \rangle \mid B_{12}\langle A!req; \dots \rangle \mid (AB_{12}):[] \mid (B_{12}A):[] \end{array}$$

Choreography Synthesis as Compliance. We briefly introduce the compliance relation that tells whether some contracts can be combined to describe a correct interaction. We reuse the main results from [14]: a typing system which assigns a (unique) global type to a set of local types. We say that a set of contracts (i.e., local types) is *compliant* if it can be assigned a choreography, i.e., a global type.

For simplicity, we use only a subset of the global types originally supported (we conjecture that extending this would not pose any difficulties). The main difference is that, in the style of [8, 10], we replace channels with participant names.

The syntax of global types is as follows:

$$\mathcal{G} ::= A \rightarrow B:e;\mathcal{G} \mid \mathcal{G} + \mathcal{G}' \mid \mathcal{G} \mid \mathcal{G}' \mid \mu \mathbf{x}.\mathcal{G} \mid \mathbf{x} \mid \mathbf{0}$$

where the first production means that a participant A sends a message of sort e to B, then interactions in \mathcal{G} take place; $\mathcal{G} + \mathcal{G}'$ means that either interactions in \mathcal{G} , or in \mathcal{G}' take place; $\mathcal{G} \mid \mathcal{G}'$ means that interactions in \mathcal{G} and \mathcal{G}' are executed concurrently; the rest of the productions are for recursive interactions, and end.

Similarly to the original synthesis, we use judgements of the form $\Gamma \vdash T \blacktriangleright \mathcal{G}$, where Γ is an environment to keep track of recursion variables, T is a system of contracts, and \mathcal{G} is the global type assigned to T . We say that a system of contracts T has global type \mathcal{G} , if one can infer the judgement $\circ \vdash T \blacktriangleright \mathcal{G}$ from the rules in the online appendix [13] (simplified from [14]) where \circ is the empty context Γ . Essentially, the synthesis rules allow to execute a set of contracts step-by-step, while keeping track of the structure of the interactions in a global type.

The main properties that we are interested in — and that are guaranteed by the synthesis — is that the inferred global type is (i) well-formed, and (ii) projectable back to the original contracts. Essentially, this means that each local type must be single-threaded, and that *knowledge of choice* is preserved — i.e., each choice is made by exactly one participant, and all the others are either made aware of the choice, or they have the same behaviour whatever choice is made.

Example 2.2. Building up on the example from Section 1.2, we combine the contract of store A with those of customers B₁ and B₂, and we obtain the system:

$$\begin{aligned}
T_{AB_1B_2} &= A\langle c_A \{B_1/b_1\} \{B_2/b_2\} \rangle \mid B_1\langle c_{B_1} \{A/a\} \{B_2/b'_2\} \rangle \mid B_2\langle c_{B_2} \{A/a'\} \{B_1/b'_1\} \rangle \\
&= A\langle B_1?req; B_2?req; B_1!quote; (B_1?order; B_2!ok + B_1?bye; B_2!bye) \rangle \\
&\quad \mid B_1\langle A!req; A?quote; (B_2!ok; A!order \oplus B_2!bye; A!bye) \rangle \\
&\quad \mid B_2\langle A!req; (B_1?ok; A?ok + B_1?bye; A?bye) \rangle
\end{aligned}$$

which can be assigned the following global type:

$$\begin{aligned}
\mathcal{G}_{AB_1B_2} &= B_1 \rightarrow A : req; B_2 \rightarrow A : req; A \rightarrow B_1 : quote; \\
&\quad (B_1 \rightarrow B_2 : ok; B_1 \rightarrow A : order; A \rightarrow B_2 : ok \quad + \quad B_1 \rightarrow B_2 : bye; B_1 \rightarrow A : bye; A \rightarrow B_2 : bye)
\end{aligned}$$

that is to say that $\circ \vdash T_{AB_1B_2} \blacktriangleright \mathcal{G}_{AB_1B_2}$ holds. Instead, if we combine the store A with B₁₂ we have

$$\begin{aligned}
T_{AB_{12}} &= A\langle c_A \{B_{12}/b_1, b_2\} \rangle \mid B_{12}\langle c_{B_{12}} \{A/a''\} \rangle \\
&= A\langle B_{12}?req; B_{12}?req; B_{12}!quote; (B_{12}?order; B_{12}!ok + B_{12}?bye; B_{12}!bye) \rangle \\
&\quad \mid B_{12}\langle A!req; A!req; A?quote; A!order; A?ok \rangle \\
\mathcal{G}_{AB_{12}} &= B_{12} \rightarrow A : req; B_{12} \rightarrow A : req; A \rightarrow B_{12} : quote; B_{12} \rightarrow A : order; A \rightarrow B_{12} : ok
\end{aligned}$$

and, again, the judgement $\circ \vdash T_{AB_{12}} \blacktriangleright \mathcal{G}_{AB_{12}}$ holds.

3 A Multiparty Version of CO₂

We introduce a version of the CO₂ calculus (for COntract-Oriented computing) [4] adapted to multiparty contracts and sessions. Let \mathbb{S} and \mathcal{S} be disjoint sets of, respectively, *session names* (ranged over by s, s', \dots) and *session variables* (ranged over by x, y, z, \dots). Let u, v, \dots range over $\mathbb{S} \cup \mathcal{S}$.

Syntax & Semantics. The syntax of CO₂ is given by the following productions:

$$\begin{array}{ll}
\text{Processes} & P, Q ::= \sum_{i \in I} P_i \cdot P_i \mid P \mid Q \mid (\vec{u}, \vec{a})P \mid X(\vec{u}, \vec{a}) \mid \mathbf{0} \\
\text{Prefixes} & p ::= \tau \mid \text{tell}_{\vec{a} \downarrow u} c \mid \text{fuse} \mid \text{do}_{\vec{a}}^u e \\
\text{Latent contracts} & K ::= \downarrow_u A \text{ says } c \mid K \mid K \\
\text{Systems} & S ::= A[P] \mid A[K] \mid s[T] \mid S \mid S \mid (\vec{u}, \vec{a})S \mid \mathbf{0}
\end{array}$$

CO₂ features CCS-style processes, equipped with branching \sum (not to be confused with the choice operator used in contracts), parallel composition \mid , restrictions of session and participant variables, and named process invocation. The prefixes are for internal action (τ), contract advertisement ($\text{tell}_{\vec{a} \downarrow u}$), session creation upon contractual agreement (fuse), and execution of contractual actions (do). A latent contract of the form $\downarrow_u A \text{ says } c$ represents the promise of participant A to fulfil c by executing do -actions on a session variable u . CO₂ systems may be parallel compositions of processes $A[P]$ (where A is the participant executing P), *latent* contracts $A[K]$ (where A is the participant to which the contracts in K have been advertised), and established sessions $s[T]$ (where s is a session name, and T is a system of *stipulated* contracts as in Section 2). We assume well-formed systems where each participant A has at most one process $A[P]$. Note that CO₂ process and system productions allow to delimit both session names/variables (\vec{u}) and participant variables (\vec{a}), but *not* participant names, which are considered public.

exists. The first two conditions, on σ and π , guarantee that all the session and participant variables are indeed instantiated. The third condition ensures that within a contract c_i , belonging to A_i , no free participant variable in c_i is substituted by A_i itself. Note that due to the condition imposed on K , each participant may have at most one contract per session.

Example 3.3. We now illustrate how Definition 3.2 works. Consider the following CO₂ system, with A , B_1 , B_2 from Section 1.2, and $T_{AB_1B_2}$ from Example 2.2:

$$\begin{aligned} S_1 &= (x, y, z) (A[\text{tell}_A \downarrow_x c_A \cdot \text{fuse} \cdot P_A] \mid B_1[\text{tell}_A \downarrow_y c_{B_1} \cdot P_{B_1}] \mid B_2[\text{tell}_A \downarrow_z c_{B_2} \cdot P_{B_2}]) \\ \rightarrow \rightarrow \rightarrow & (x, y, z) (A[\text{fuse} \cdot P_A] \mid A[\downarrow_x A \text{ says } c_A \mid \downarrow_y B_1 \text{ says } c_{B_1} \mid \downarrow_z B_2 \text{ says } c_{B_2}] \mid B_1[P_{B_1}] \mid B_2[P_{B_2}]) \\ \xrightarrow{A: \text{fuse}} (s) S'_1 &= (s) (A[P_A] \sigma \pi \mid s[T_{AB_1B_2} \mid Q(T_{AB_1B_2})] \mid B_1[P_{B_1}] \sigma \pi \mid B_2[P_{B_2}] \sigma \pi) \\ &\text{where } \sigma = \{s/x, y, z\} \text{ and } \pi = \{A/a, a', B_1/b_1, b'_1, B_2/b_2, b'_2\} \end{aligned}$$

The initial system S_1 is the one considered in Section 1.2, where all the participants are ready to advertise their respective contracts to the store A , by using a $\text{tell}_A \downarrow$ -primitive. This has the effect of creating corresponding latent contracts within A . Once all the latent contracts are in a same location, they may be fused. In this case, given σ and π as above, Definition 3.2 is indeed applicable: the domains of σ and π comply with the definition's premises, and we already saw that a system consisting of c_A , c_{B_1} and c_{B_2} may be assigned a global type. Hence, a new session s is created, based on the system of contracts $T_{AB_1B_2}$, plus the queues connecting all pairs of participants. The session variables of the latent contracts being fused (i.e., x for participant A , y for B_1 , and z for B_2) are all substituted with the fresh session name s in the processes P_A , P_{B_1} and P_{B_2} , via σ . Similarly for participant variables which are substituted with participant names, via π .

The CO₂ semantic rules are to be considered up-to a standard structural congruence relation \equiv (cf. [13]): we just point out that $A[K] \mid A[K'] \equiv A[K \mid K']$ allows to select a compliant subset from a group of latent contracts, before performing a fuse — thus adding flexibility to the synthesis of choreographies.

Example 3.4. Consider the system:

$$\dots B[\text{fuse} \cdot P \mid Q] \mid B[\downarrow_x A_1 \text{ says } a!\text{int} \mid \downarrow_y A_2 \text{ says } a'?\text{int} \mid \downarrow_z A_3 \text{ says } b?\text{bool}] \dots$$

The fuse prefix cannot be fired: no contract matches A_3 's, and thus the three latent contracts cannot be assigned a global type. However, by rearranging the system with congruence \equiv , we have:

$$\dots B[\text{fuse} \cdot P \mid Q] \mid B[\downarrow_x A_1 \text{ says } a!\text{int} \mid \downarrow_y A_2 \text{ says } a'?\text{int}] \mid B[\downarrow_z A_3 \text{ says } b?\text{bool}] \dots$$

It is now possible to synthesise a global type $A_1 \rightarrow A_2 : \text{int}$, and a session may be created for A_1 and A_2 . A_3 's latent contract may be fused later on.

3.1 Flexibility of Session Establishment

We highlight the flexibility of our definition of contract agreement, together with the semantics of CO₂, by discussing three features: (i) contracts may use a mix of participant names and variables, (ii) different contracts may use common participant variables, and (iii) the definition of agreement may be easily extended.

Contracts with both participant names and variables. A may want to sell an item to a *specific* participant B, via *any* shipping company that provides a package tracking system. A’s contract may be:

`B!price;B?ack;a!request;a?tracking;B!tracking`

saying that the seller A must send a price to the buyer B; once B has acknowledged, A must send a shipping request to a shipper a — who must send back a tracking number, which is then forwarded to B. This contract may be fused only if B takes part in the session, while the role of shipper a may be played by any participant.

Contracts sharing participant variables. Consider the CO₂ process:

$\dots A[\text{tell}_{A \downarrow x}(\mathbf{b}!request) \dots X(\vec{z}, \mathbf{b})] \dots$ where $X(\vec{z}, \mathbf{b}) := (y, \mathbf{b}') \text{tell}_{A \downarrow y}(\mathbf{b}'!quote; \dots \mathbf{b}!address) \dots X(\vec{z}, \mathbf{b})$

Here, A advertises two contracts: the first one (`b!request`) is used by A to find a shipping company, and the second (`b'!quote; ... b!address`) to sell items. The two contracts are linked by the common variable `b`: whenever the first one is fused, variable `b` is instantiated to a participant name, say B, which is also substituted in the second. This means that whenever a new selling session starts, B will also be involved as the receiver of the address message.

Possible extensions. The participants firing fuse-primitives are playing the role of brokers in our framework. Depending on their implementation, brokers may also have some obligations in the contracts they fuse, or they may want to enforce some general policy — therefore they may have additional requirements before agreeing to start a session. For instance, a broker may not want to start a session with too many participants as it may be too resource demanding (too many connections etc.). Another broker may want to start sessions that terminates after a limited number of interactions, because it has a short life expectancy, e.g., due to an approaching scheduled maintenance. Another kind of broker may precisely want to start sessions which do not terminate, e.g., if the broker is interested in resilient services.

Several variations of the fuse primitive are possible thanks to the fact that we base contract agreements on objects representing the overall choreography. We introduce `fuse[n]`, a version of fuse that only fuses sessions where there are at least n participants; `fuseT`, which has the additional constraint that no recursive behaviour is allowed in the synthesised choreography (therefore ensuring that the session will eventually terminate), and `fuseR`, which only creates sessions when the synthesised choreography never terminates (i.e., it only consists of recursive behaviours).

The three extensions may be defined directly via small modifications of Definition 3.2:

- `fuse[n]`: we add the condition $|\mathcal{P}(\mathcal{G})| \geq n$, where $\mathcal{P}(\mathcal{G})$ is the set of participants in \mathcal{G} ;
- `fuseT`: we add the condition that there should not be any recursion variable `x` in \mathcal{G} ;
- `fuseR`: we add the condition that `0` does not appear in \mathcal{G} .

This kind of properties must be checked for at global level because it cannot always be decided by looking at the individual contracts. For instance, a participant might exhibit a recursive behaviour in one of the branches of an external choice, while the participant it interacts with may always choose a branch that is not recursive. Note that none of these variations actually affect the results that follow, since the original fuse primitive is also blocking. The variations only restrict some of its applications. Further variations of fuse are sketched in Section 6, as future work.

4 The Problem of Honesty

In this section, we discuss and define the notion of *honesty* [4], i.e., the ability of a participant to always fulfil its contracts, in any context. Broadly speaking, in our contract-oriented setting, honesty is the counterpart of well-typedness in a session type setting: the static proof that a participant always honours its contracts provides guarantees about its runtime behaviour.

As seen in Example 3.1, each do prefix within the process of a participant, say $A[P]$, is driven by the contract that A promised to abide by. In a sense, CO_2 is *culpability-driven*, according to Definition 4.1 below: when a participant is culpable, it has the duty of making the session progress according to its contract.

Definition 4.1 (Culpability). *Let S be a CO_2 system with a session s , i.e., $S \equiv (\vec{u}, \vec{a}) (A[P] \mid s[T] \mid S_1)$. We say that A is culpable in S (at session s) when there exist B and e such that $T \xrightarrow{A \Leftarrow B; e}$.*

A culpable participant can overcome its status by firing its do prefixes, according to $[\text{CO}_2\text{-Do}]$, until another participant becomes culpable or session s terminates. Hence, as long as a culpable participant A does not enable a do-prefix matching a contractual action, A will remain culpable. Note that when a participant is involved in multiple sessions, it may result culpable in more than one of them.

When a participant A is always able to fulfil its contractual actions (i.e., overcome its culpability), no matter what other participants do, then it is said to be *honest* (cf. Definition 4.9). This is a desirable property in a distributed contract-oriented scenario: a participant may be stuck in a culpable condition either due to “simple” bugs (cf. Example 4.7), or due to the unexpected (or malicious) behaviour of other participants (cf. Example 5.6). Therefore, before deploying a service, its developers might want to ensure that it will always be able exculpate itself.

Formally, as in [1], we base the definition of honesty on the relationship between the ready sets of a contract, and those of a CO_2 process. We call the former *contract ready sets*, and the latter *process ready sets*. The concept of contract ready sets is similar to [9, 4, 1], where only bilateral contracts are considered. Here, we adapt it to suit our multiparty contract model.

Definition 4.2 (Contract Ready Sets). *The ready sets of a contract c , written $\text{CRS}(c)$, are:*

$$\text{CRS}(c) = \begin{cases} \text{CRS}(c') & \text{if } c = \mu x.c' \\ \{\{(A_i, e_i) \mid i \in I\}\} & \text{if } c = \bigoplus_{i \in I} A_i!e_i; c_i \text{ and } I \neq \emptyset \\ \{\{(A, e_i) \mid i \in I\}\} & \text{if } c = \sum_{i \in I} A?e_i; c_i \end{cases}$$

Intuitively, when a participant A is bound to a contract c , the ready sets of c tell which interactions A must be able to perform towards other participants. Each interaction has the form of a pair, consisting of a participant name and a message sort. The interactions offered by an external choice are all available at once, while those offered by an internal choice are mutually exclusive.

Example 4.3. *Consider the system of contracts $T_{AB_1B_2}$ from Example 2.2 — and in particular, the stipulated contracts therein, with substitution $\pi = \{A/a, a', B_1/b_1, b'_1, B_2/b_2, b'_2\}$ from Example 3.3:*

$$\begin{aligned} \tilde{c}_A &= c_A \pi = B_1?req; B_2?req; B_1!quote; (B_1?order; B_2!ok + B_1?bye; B_2!bye) \\ \tilde{c}_{B_1} &= c_{B_1} \pi = A!req; A?quote; (B_2!ok; A!order \oplus B_2!bye; A!bye) \\ \tilde{c}_{B_2} &= c_{B_2} \pi = A!req; (B_1?ok; A?ok + B_1?bye; A?bye) \end{aligned}$$

We have $\text{CRS}(\tilde{c}_A) = \{\{(B_1, req)\}\}$: in other words, at this point of the contract, an interaction is expected between A and B_1 (since A is waiting for req), while no interaction is expected between A and B_2 .

Let us now equip $T_{AB_1B_2}$ with one queue between each pair of participants, and let it perform the request exchange between B_1 and A , with the transitions:

$$T_{AB_1B_2} \mid Q(T_{AB_1B_2}) \xrightarrow{B_1 \rightarrow A:\text{req}} \xrightarrow{A \leftarrow B_1:\text{req}} T'_{AB_1B_2} \mid Q(T_{AB_1B_2})$$

We have that \tilde{c}_A in $T'_{AB_1B_2}$ is now reduced to:

$$\tilde{c}_A' = B_2?\text{req}; B_1!\text{quote}; (B_1?\text{order}; B_2!\text{ok} + B_1?\text{bye}; B_2!\text{bye})$$

and thus we have $\text{CRS}(\tilde{c}_A') = \{\{(B_2, \text{req})\}\}$, i.e., A is now waiting for a request from B_2 .

If we let the system reduce further, \tilde{c}_A' reaches its external choice:

$$\tilde{c}_A'' = B_1?\text{order}; B_2!\text{ok} + B_1?\text{bye}; B_2!\text{bye}$$

Now, the ready sets become $\text{CRS}(\tilde{c}_A'') = \{\{(B_1, \text{order}), (B_1, \text{bye})\}\}$, i.e., A must handle both answers from B_1 . Instead, when \tilde{c}_{B_1} reduces to its internal choice, we have:

$$\tilde{c}_{B_1}'' = B_2!\text{ok}; A!\text{order} \oplus B_2!\text{bye}; A!\text{bye}$$

Thus, its ready sets become $\text{CRS}(\tilde{c}_{B_1}'') = \{\{(B_2, \text{ok})\}, \{(B_2, \text{bye})\}\}$: B_1 is free to choose either branch.

Example 4.3 shows that, when a contract c of a principal A evolves within a system T , its ready sets change. Now we need to define the counterpart of contract ready sets for CO_2 processes, i.e., the *process ready sets*. Again, we adapt the definition from [1] to our multiparty contract model.

Definition 4.4 (Process Ready Set). *For all CO_2 systems S , all participants A, B and sessions u , we define the set of pairs:¹*

$$\text{PRS}_A^u(S) = \{(B, e) \mid \exists \vec{v}, \vec{a}, P, P', Q, S' : S \equiv (\vec{v}, \vec{a}) (A[\text{do}_B^u e. P + P' \mid Q] \mid S_0) \mid S_1 \wedge u \notin \vec{v}\}$$

Intuitively, Definition 4.4 says that the process ready set of A over a session u in a system S contains the interactions that A is immediately able to perform with other participants through its do_B^u prefixes. Just as in contract ready sets, the interactions are represented by participant/sort pairs.

Next, we want to characterise a weaker notion of the process ready set, so it only takes into account the first actions *on a specific session* that a participant is ready to make.

Definition 4.5 (Weak Process Ready Set). *We write $S \xrightarrow{\neq(A: \text{do}^u)} S'$ iff:*

$$\exists B, p: S \xrightarrow{B:p} S' \implies (A \neq B \vee \forall e: \forall C: p = \text{do}_C^v e \implies u \neq v)$$

We then define the set of pairs $\text{WPRS}_A^u(S)$ as:

$$\text{WPRS}_A^u(S) = \left\{ (B, e) \mid \exists S': S \xrightarrow{\neq(A: \text{do}^u)} S' \wedge (B, e) \in \text{PRS}_A^u(S') \right\}$$

In Definition 4.5, we are not interested in the actions that do not relate to the session u . Thus, we allow the system to evolve either by (i) letting any other participant other than A do an action, or (ii) letting A act on a different session than u , or (iii) do internal actions.

We now introduce the final ingredient for honesty, that is the notion of *readiness* of a participant.

¹The side condition “ $u \notin \vec{v}$ ” of Definition 4.4 deals with cases like $S_0 = (s) (A[\text{do}_B^s \text{int}])$ and $S = S_0 \mid s[A\langle B!\text{int} \rangle \mid \dots] \mid \dots$: without the side condition, $\text{PRS}_A^u(S_0) = \{\{(B, \text{int})\}\}$ — hence, by Def. 4.6, A would result to be ready in S .

Definition 4.6 (Readiness). We say that A is ready in S iff, whenever $S \equiv (\vec{u}, \vec{b})S_0$ for some \vec{u}, \vec{b} and $S_0 \equiv s[A\langle c \rangle \mid \dots] \mid \dots$, the following holds:

$$\exists \mathcal{X} \in \text{CRS}(c) : ((B, e) \in \mathcal{X} \implies (B, e) \in \text{WPRS}_A^s(S_1))$$

Definition 4.6 says that a participant A is *ready* in a system S whenever its process ready sets for a session s will eventually contain all the participant/sort pairs of one of the contract ready sets of A 's contract in s . When a participant A is “ready”, then, for any of its contracts c , the CO_2 process of A is (eventually) able to fulfil at least the interactions in c 's prefix.

Example 4.7. We have seen that, after fusion of the latent contracts of S_1 (in Ex. 3.3) we obtain:

$$(s)S'_1 \equiv (s)(A[P_A\sigma\pi] \mid s[T_{AB_1B_2} \mid Q(T_{AB_1B_2})] \mid B_1[P_{B_1}\sigma\pi] \mid B_2[P_{B_2}\sigma\pi])$$

where the substitutions σ and π are also from Ex. 3.3. Let us define the processes (after substitutions):

$$\begin{aligned} P_A\sigma\pi &= \text{do}_{B_1}^s \text{req} . \text{do}_{B_2}^s \text{req} . \text{do}_{B_1}^s \text{quote} . (\text{do}_{B_1}^s \text{order} . \text{do}_{B_2}^s \text{ok} + \text{do}_{B_1}^s \text{bye} . \text{do}_{B_2}^s \text{bye}) \\ P_{B_1}\sigma\pi &= \tau . \text{do}_A^s \text{req} . \text{do}_A^s \text{quote} . \text{do}_A^s \text{order} \\ P_{B_2}\sigma\pi &= \text{do}_A^s \text{req} . (\text{do}_{B_1}^s \text{ok} . \text{do}_A^s \text{ok} + \text{do}_{B_1}^s \text{bye} . \text{do}_A^s \text{bye}) \end{aligned}$$

Thus, we have:

$$\begin{aligned} \text{PRS}_A^s(S'_1) &= \{(B_1, \text{req})\} = \text{WPRS}_A^s(S'_1) \\ \text{PRS}_{B_1}^s(S'_1) &= \emptyset \neq \{(A, \text{req})\} = \text{WPRS}_{B_1}^s(S'_1) \\ \text{PRS}_{B_2}^s(S'_1) &= \{(A, \text{req})\} = \text{WPRS}_{B_2}^s(S'_1) \end{aligned}$$

Note that the τ prefix in P_{B_1} prevents B_1 from interacting immediately with A on session s , although it is “weakly ready” to do it. Hence, considering that the weak process ready sets of each participant in S'_1 match their respective contract ready sets in $T_{AB_1B_2}$ (Example 4.3) according to Definition 4.6 we have that participants A , B_1 and B_2 are all ready in $(s)S'_1$.

Before defining honesty formally, we need to characterise the class of systems for which this concept is meaningful, i.e., those systems where a participant is not (yet) involved in latent contracts nor active sessions.

Definition 4.8 (Initial System). A CO_2 system S is A -initial if S has no sub-term of the form $\downarrow_u A \text{ says } c$ or $A\langle c \rangle$ with $c \neq \mathbf{0}$. A CO_2 system S is initial when it is A -initial for each participant A in S .

Definition 4.9 (Honesty). We say that $A[P]$ is honest iff, for all A -initial $S \equiv (\vec{u}, \vec{b})(A[P] \mid S_0)$ s.t. $S \rightarrow^* S'$, A is ready in S' .

A process $A[P]$ is said to be honest when, for all contexts and reductions that $A[P]$ may be engaged in, A is persistently ready. In other words, there is a continuous correspondence between the interactions exposed in the contract ready sets and the process ready sets of the possible reductions of any system involving $A[P]$. The definition rules out contexts with latent/stipulated contracts of A , otherwise A could be made trivially dishonest, e.g. by inserting a latent contract $\downarrow_u A \text{ says } c$ that A cannot fulfil.

Example 4.10. Consider the process $B_1[\text{tell}_A \downarrow_y c_{B_1} . P_{B_1}]$ of system S_1 , as defined in Examples 3.3 and 4.7. We show that this process is not honest. In fact, S_1 can reduce as $S_1 \rightarrow^* (s)S'_1 \rightarrow^* (s)S''_1$, where:

$$\begin{aligned} (s)S''_1 &= (s) \left(A[\text{do}_{B_1}^s \text{order} . \text{do}_{B_2}^s \text{ok} + \text{do}_{B_1}^s \text{bye} . \text{do}_{B_2}^s \text{bye}] \right. \\ &\quad \mid s[A\langle B_1?order; B_2!ok + B_1?bye; B_2!bye \rangle \\ &\quad \quad \mid B_1\langle B_2!ok; A!order \oplus B_2!bye; A!bye \rangle \mid B_2\langle B_1?ok; A?ok + B_1?bye; A?bye \rangle \\ &\quad \quad \mid (AB_1):[] \mid (B_1A):[] \mid (AB_2):[] \mid (B_2A):[] \mid (B_1B_2):[] \mid (B_2B_1):[]] \\ &\quad \left. \mid B_1[\text{do}_A^s \text{order}] \mid B_2[\text{do}_{B_1}^s \text{ok} . \text{do}_A^s \text{ok} + \text{do}_{B_1}^s \text{bye} . \text{do}_A^s \text{bye}] \right) \end{aligned}$$

At this point, we see that there is a problem in the implementation of B_1 : it does not notify the other buyer before making an order. In fact, B_1 's process is trying to perform $\text{do}_A^s \text{order}$, but its contract requires that $\text{do}_{B_2}^s \text{ok}$ is performed first (or $\text{do}_{B_2}^s \text{bye}$, if the quote is rejected). This is reflected by the mismatch between B_1 's process ready set in S_1'' , and its contract ready sets, in session s :

$$\begin{aligned} \text{PRS}_{B_1}^s(S_1'') &= \{(A, \text{order})\} \\ \text{CRS}(B_2! \text{ok}; A! \text{order} \oplus B_2! \text{bye}; A! \text{bye}) &= \{(B_2, \text{ok})\}, \{(B_2, \text{bye})\} \end{aligned}$$

In terms of the above definitions, there exists a system S_1 — containing $B_1[\text{tell}_A \downarrow_y c_{B_1} . P_{B_1}]$ — that reduces to a $(s)S_1''$ where B_1 is not ready (Definition 4.6). Therefore, $B_1[\text{tell}_A \downarrow_y c_{B_1} . P_{B_1}]$ is not honest. In fact, B_1 is culpable in $(s)S_1''$, according to Definition 4.1.

As in [1], the definition of honesty subsumes a *fair* scheduler, eventually allowing participants to fire persistently (weakly) enabled do actions.

Honesty is not decidable in general [4], but for a bilateral contract model it has been approximated either via an abstract semantics [4] or a type discipline [1] for CO_2 . We believe that these approximations may be easily adapted to our setting (see Section 6 for more details).

5 Results

We now give the main properties of our framework. We ensure that two basic features of CO_2 hold in our multiparty adaptation: the state of a session always allows to establish who is responsible for making the system progress (Th. 5.1) and honest participants can always exculpate themselves (Th. 5.3). We then formalise a link between the honesty of participants, and two key properties borrowed from the session types setting: Th. 5.4 introduces session fidelity in CO_2 ; and Th. 5.5 introduces a notion of progress in CO_2 , based on the progress of the contractual agreement (and its choreography).

Theorem 5.1 (Unambiguous Culpability). *Given an initial CO_2 system S , if $S \rightarrow^* S' \equiv (\vec{u}, \vec{b})(s[T] \mid \dots)$ such that $T \not\equiv \mathbf{0}$, then there exists at least one culpable participant in S' .*

Theorem 5.1 says that in an active session established through a fuse reduction, there is always at least one participant $A[P]$ who leads the next interaction. Thus, if a corresponding $\text{do}_B^s e$ prefix is not in P , S may get stuck, and A is culpable.

Example 5.2. *Consider the system S_1'' in Example 4.10, and the system of contracts in its session s :*

$$\begin{aligned} T_s &= A\langle B_1? \text{order}; B_2! \text{ok} + B_1? \text{bye}; B_2! \text{bye} \rangle \\ &\quad | B_1\langle B_2! \text{ok}; A! \text{order} \oplus B_2! \text{bye}; A! \text{bye} \rangle | B_2\langle B_1? \text{ok}; A? \text{ok} + B_1? \text{bye}; A? \text{bye} \rangle \\ &\quad | (AB_1):[] | (B_1A):[] | (AB_2):[] | (B_2A):[] | (B_1B_2):[] | (B_2B_1):[] \end{aligned}$$

We have $T_s \xrightarrow{B_1 \Leftarrow B_2: \text{ok}}$ and $T_s \xrightarrow{B_1 \Leftarrow B_2: \text{bye}}$. Hence, B_1 is responsible for the next interaction, and culpable for S_1'' being stuck.

Theorem 5.3 (Exculpation). *Given an A -initial CO_2 system S_0 with $A[P]$ honest, whenever $S_0 \rightarrow^* S \equiv (\vec{u}, \vec{a})(s[T] \mid S_1)$ and A is culpable in S at session s , there exist B and e such that: $S \xrightarrow{A: P}^* \xrightarrow{A: \text{do}_B^s e}$ where $p = \tau$ or $p = \text{tell}_\downarrow _$.*

Theorem 5.3 follows from the definition of honesty, formalising that honest participants can always overcome their culpability, by firing their contractual do actions (possibly after advertising other contracts or performing some internal actions).

Theorem 5.4 (Fidelity). *For all initial systems S with only honest participants, if S is such that $S \rightarrow^* S' \equiv (\vec{u}, \vec{a}) (A[P] \mid s[T] \mid S_0)$, then $(S' \xrightarrow{\neq(A: \text{do}^s)} \xrightarrow{A: \text{do}_B^s e}) \iff (T \xrightarrow{A \equiv B: e})$ (where $\xrightarrow{\neq(A: \text{do}^u)}$, as in Def. 4.5, intuitively stands for any reduction not involving session s).*

Theorem 5.4 says that each (honest) participant will strictly adhere to its contracts, once they have been fused in a session. It follows directly from the semantics of CO_2 (that forbid non-contractual do prefixes to be fired) and from the definition of honesty.

Theorem 5.5 below introduces the notion of global progress, which is slightly different from the contractual progress. In fact, progress in CO_2 is only meaningful *after* a session has been established, and thus a culpable participant exists. A system without sessions may not progress because a set of compliant contracts cannot be found, or a fuse prefix is not enabled. In both cases, no participant may be deemed culpable, and thus responsible for the next move. However, the system may progress again if other (honest) participants join it, allowing a session to be established.

Theorem 5.5 (Global Progress). *Given an initial CO_2 system S_0 with only honest participants, if $S_0 \rightarrow^* S \equiv (\vec{u}, \vec{a}) (s[T] \mid S_1)$ with $T \not\equiv \mathbf{0}$, then $S \rightarrow$.*

Theorem 5.5 follows from the definition of honesty (i.e., participants are always ready to fulfil their contracts), the fact that contract compliance guarantees contractual progress [14], Theorem 5.3, and the semantics of CO_2 (in particular, rule $[\text{CO}_2\text{-Do}]$). This result also holds for systems where a process takes part in multiple sessions: the honesty of all participants guarantees that all sessions will be completed.

Example 5.6. *We now give a simple example on a system with multiple sessions. We show how a seemingly honest process (B) could be deemed culpable due to the unexpected behaviour of other participants, and how honest participants guarantee progress of the whole system. Consider:*

$$S = (x, y, z, w) \left(A[\text{tell}_A \downarrow_x (B!\text{int}) . \text{fuse} . \text{fuse}] \mid B[\text{tell}_A \downarrow_y (A?\text{int}) . \text{tell}_A \downarrow_z (C!\text{bool}) . \text{do}_A^y \text{int} . \text{do}_C^z \text{bool}] \right. \\ \left. \mid C[\text{tell}_A \downarrow_w (B?\text{bool}) . \text{do}_B^w \text{bool}] \right)$$

After all four contracts have been advertised to A and fused, the system reduces to:

$$S' = (s_1, s_2) \left(A[\mathbf{0}] \mid B[\text{do}_A^{s_1} \text{int} . \text{do}_C^{s_2} \text{bool}] \mid C[\text{do}_B^{s_2} \text{bool}] \right. \\ \left. \mid s_1[A\langle B?\text{int} \rangle \mid B\langle A!\text{int} \rangle \mid (AB): [] \mid (BA): []] \mid s_2[B\langle C!\text{bool} \rangle \mid C\langle B?\text{bool} \rangle \mid (BC): [] \mid (CB): []] \right)$$

Even if both sessions s_1 and s_2 enjoy contractual progress, S' is stuck: A does not perform the promised action, thus remaining culpable in s_1 ; B is stuck waiting in s_1 , thus remaining culpable in s_2 .² Indeed, neither A nor B are ready in S' , and thus they are not honest in S . Hence, global progress is not guaranteed. Let us now consider the following variant of S , where all participants are honest:

$$\hat{S} = (x, y, z, w) \left(A[(\text{tell}_A \downarrow_x (B!\text{int}) . \text{do}_B^x \text{int}) \mid \text{fuse} \mid \text{fuse}] \mid C[\text{tell}_A \downarrow_w (B?\text{bool}) . \text{do}_B^w \text{bool}] \right. \\ \left. \mid B[\text{tell}_A \downarrow_y (A?\text{int}) . \text{tell}_A \downarrow_z (C!\text{bool}) . (\text{do}_A^y \text{int} . \text{do}_C^z \text{bool} + \tau . (\text{do}_A^y \text{int} \mid \text{do}_C^z \text{bool}))] \right)$$

In this case, A will respect its contractual duties, while B will be ready to fulfil its contracts on both sessions — even if one is not activated, or remains stuck (here, τ represents an internal action, e.g., a timeout: if the first $\text{do}_A^y \text{int}$ cannot reduce, B falls back to running the sessions in parallel). The honesty of all participants in \hat{S} guarantees that, once a session is active, it will reach its completion.

²In this case, B is deemed culpable in s_2 because its implementation did not expect A to misbehave.

6 Concluding remarks

In this work, we investigated the combination of the contract-oriented calculus CO_2 with a contract model that fulfils two basic design requirements: (i) it supports multiparty agreements, and (ii) it provides an explicit description of the choreography that embodies each agreement. These requirements prompted us towards the well-established results from the session types setting — in particular, regarding the interplay between a well-formed global type and its corresponding local behaviours. We introduced the concepts of global progress and session fidelity in CO_2 , also inspired from the analogous concepts in theories based on session types. We built our framework upon a simple version of session types, and yet it turns out to be quite flexible, allowing for sessions where the number of participants is not known beforehand.

Related work. The origin of CO_2 goes back to [5]. The calculus was generalised in [2, 3] to suit different contract models (e.g., contracts as processes or logic formulæ). In [4], it has been instantiated to a theory of bilateral contracts inspired by [9]. A negative result in [4] is that the problem of honesty (Section 4) is not decidable. A type system for CO_2 processes providing a decidable approximation of honesty was introduced in [1]. This result relies on the product between a finite state system (approximating contracts) and a Basic Parallel Process (approximating a CO_2 process). Considering that the systems of contracts in this work form a (strict) subset of the local/global types in [10], for which each configuration is reachable by a 1-buffer execution, we believe that the type system in [1] may be adapted to our setting.

The seminal top-down approach of multiparty session types has been first described in [11]. In summary, the framework works as follows: designers specify a choreography (i.e., a global type), which is then projected onto local behaviours (i.e., local types), which in turn are used to type-check processes. A dual approach was introduced in [14]: from a set of local types it is possible to synthesise a choreography (i.e., a global type). This is precisely the result we use as a basis for our definition of compliant contracts.

The semantic correspondence between global types and projected local behaviours has been investigated in [12, 8]. To the best of our knowledge, no other contract model besides ours is based on explicit choreography synthesis. A related approach is presented in [7]: multiple contracts are considered compliant when their composition (i.e., the system of contracts) guarantees completion. In our work, progress (subsuming completion) is provided by the synthesis of a global type. In [6], contracts are considered compliant when their composition adheres to a predetermined choreography; in our framework, however, no choreography is assumed beforehand.

Future work. We plan to extend our work so to offer even more flexibility. For example, by introducing a parameterised fuse primitive which starts a session according to different criteria, when more than one agreement is possible (as in our introductory example). For instance, one could choose the agreement involving the most (or least) number of participants. These criteria may be based on a semantic characterisation of global types, e.g., as the ones in [12, 8]. We also plan to study the possibility for a participant to be involved in a session under multiple contracts, e.g., a bank advertising two services, and a customer publishing a contract which uses both of them in a well-formed choreography.

Another research direction is the concept of “group honesty”. In fact, the current definition of honesty is quite strict: it basically verifies each participant in isolation, thus providing a sufficient (but not necessary) condition for progress. Consider, for example, a CO_2 system like:

$$S = (x, y) \left(A[\text{tell}_A \downarrow_x (B!int \oplus B!bool) . \text{fuse} . \text{do}_B^x \text{int}] \mid B[\text{tell}_A \downarrow_y (A?int + A?bool) . \text{do}_A^y \text{int}] \right)$$

B is dishonest, since it is not ready for the `bool` branch of its contract. However, the system S has progress: when B establishes a session with A , the latter will never take the `bool` branch; hence, B will not remain culpable. This kind of “group honesty” may be used to validate (sub-)systems of participants developed

by the same organization: it would ensure that they never “cheat each other”, and are collectively honest when deployed in any context. Furthermore, the group honesty of all participants in a system S may turn out to be a necessary condition for the global progress of S .

Acknowledgements. We would like to thank Massimo Bartoletti, Emilio Tuosto, and Roberto Zunino for their valuable advice, discussions and comments. We would also like to thank the anonymous reviewers for their suggestions. This work is partially supported by Aut. Region of Sardinia under grants L.R.7/2007 CRP-17285 (TRICS) and P.I.A. 2010 project “Social Glue”, and by MIUR PRIN 2010-11 project “Security Horizons”, and by EU COST Action IC1201 (BETTY).

References

- [1] M. Bartoletti, A. Scalas, E. Tuosto & R. Zunino (2013): *Honesty by Typing*. In: *FMOODS/FORTE*, pp. 305–320. Available at http://dx.doi.org/10.1007/978-3-642-38592-6_21.
- [2] M. Bartoletti, E. Tuosto & R. Zunino (2011): *Contracts in distributed systems*. In: *ICE*, pp. 130–147. Available at <http://dx.doi.org/10.4204/EPTCS.59.11>.
- [3] M. Bartoletti, E. Tuosto & R. Zunino (2012): *Contract-oriented Computing in CO₂*. *Scientific Annals in Comp. Sci.* 22(1), pp. 5–60. Available at <http://dx.doi.org/10.7561/SACS.2012.1.5>.
- [4] M. Bartoletti, E. Tuosto & R. Zunino (2012): *On the Realizability of Contracts in Dishonest Systems*. In: *COORDINATION*. Available at http://dx.doi.org/10.1007/978-3-642-30829-1_17.
- [5] M. Bartoletti & R. Zunino (2010): *A Calculus of Contracting Processes*. In: *LICS*. Available at <http://doi.ieeecomputersociety.org/10.1109/LICS.2010.25>.
- [6] M. Bravetti, I. Lanese & G. Zavattaro (2009): *Contract-Driven Implementation of Choreographies*. In: *TGC, LNCS 5474*. Available at http://dx.doi.org/10.1007/978-3-642-00945-7_1.
- [7] M. Bravetti & G. Zavattaro (2008): *A Foundational Theory of Contracts for Multi-party Service Composition*. *Fundam. Inform.* 89(4). Available at <http://iospress.metapress.com/content/x7613955410077k2/>.
- [8] G. Castagna, M. Dezani-Ciancaglini & L. Padovani (2012): *On Global Types and Multi-Party Session*. *Logical Methods in Comp. Sci.* 8(1). Available at [http://dx.doi.org/10.2168/LMCS-8\(1:24\)2012](http://dx.doi.org/10.2168/LMCS-8(1:24)2012).
- [9] G. Castagna, N. Gesbert & L. Padovani (2009): *A theory of contracts for Web services*. *ACM Trans. on Prog. Lang. and Sys.* 31(5). Available at <http://doi.acm.org/10.1145/1538917.1538920>.
- [10] P.M. Deniérou & N. Yoshida (2012): *Multiparty Session Types Meet Communicating Automata*. In: *ESOP*. Available at http://dx.doi.org/10.1007/978-3-642-28869-2_10.
- [11] K. Honda, N. Yoshida & M. Carbone (2008): *Multiparty asynchronous session types*. In: *POPL*. Available at <http://doi.acm.org/10.1145/1328438.1328472>.
- [12] I. Lanese, C. Guidi, F. Montesi & G. Zavattaro (2008): *Bridging the Gap between Interaction- and Process-Oriented Choreographies*. *SEFM '08*. Available at <http://dx.doi.org/10.1109/SEFM.2008.11>.
- [13] J. Lange & A. Scalas (2013): *Choreography Synthesis as Contract Agreement + Appendix*. Available at <http://tcs.unica.it/publications>.
- [14] J. Lange & E. Tuosto (2012): *Synthesising Choreographies from Local Session Types*. In: *CONCUR*. Available at http://dx.doi.org/10.1007/978-3-642-32940-1_17.
- [15] OASIS (2012): *Reference Architecture Foundation for Service Oriented Architecture*. *Comm. Spec.* 01, v.1.0. Available at <http://docs.oasis-open.org/soa-rm/soa-ra/v1.0/cs01/soa-ra-v1.0-cs01.html>.