

Modelling and Verifying Contract-Oriented Systems in Maude

Massimo Bartoletti¹, Maurizio Murgia¹, Alceste Scalas¹(✉),
and Roberto Zunino²

¹ Università Degli Studi di Cagliari, Cagliari, Italy
{bart,alceste.scalas}@unica.it, murgia88@gmail.com

² Università Degli Studi di Trento, Trento, Italy
roberto.zunino@unitn.it

Abstract. We address the problem of modelling and verifying contract-oriented systems, wherein distributed agents may advertise and stipulate contracts, but — differently from most other approaches to distributed agents — are not assumed to always behave “honestly”. We describe an executable specification in Maude of the semantics of CO₂, a calculus for contract-oriented systems [6]. The *honesty* property [5] characterises those agents which always respect their contracts, in *all* possible execution contexts. Since there is an infinite number of such contexts, honesty cannot be directly verified by model-checking the state space of an agent (indeed, honesty is an undecidable property in general [5]). The main contribution of this paper is a sound verification technique for honesty. To do that, we safely over-approximate the honesty property by abstracting from the actual contexts a process may be engaged with. Then, we develop a model-checking technique for this abstraction, we describe an implementation in Maude, and we discuss some experiments with it.

1 Introduction

Contract-oriented computing is a software design paradigm where the interaction between clients and services is disciplined through contracts [4, 6]. Contract-oriented services start their life-cycle by advertising contracts which specify their required and offered behaviour. When compliant contracts are found, a session is created among the respective services, which may then start interacting to fulfil their contracts. Differently from other design paradigms (e.g. those based on the session types discipline [10]), services are not assumed to be *honest*, in that they might not respect the promises made [5]. This may happen either unintentionally (because of errors in the service specification), or because of malicious behaviour.

Dishonest behaviour is assumed to be automatically detected and sanctioned by the service infrastructure. This gives rise to a new kind of attacks, that exploit possible discrepancies between the promised and the actual behaviour. If a service does not behave as promised, an attacker can induce it to a situation where the service is sanctioned, while the attacker is reckoned honest. A crucial

problem is then how to avoid that a service results definitively culpable of a contract violation, despite of the honest intentions of its developer.

In this paper we present an executable specification in Maude [9] of CO₂, a calculus for contract-oriented computing [4]. Furthermore, we devise and implement a sound verification technique for honesty. We start in Sect. 2 by introducing a new model for contracts. Borrowing from other approaches to behavioural contracts [5, 8], ours are bilateral contracts featuring internal/external choices, and recursion. We define and implement in Maude two crucial primitives on contracts, i.e. *compliance* and *culpability testing*, and we study some relevant properties.

In Sect. 3 we present CO₂ (instantiated with the contracts above), and an executable specification of its semantics in Maude. In Sect. 4 we formalise a weak notion of honesty, i.e. when a process P is honest *in a given context*, and we implement and experiment with it through the Maude model checker.

The main technical results follow in Sect. 5, where we deal with the problem of checking honesty in *all* possible contexts. To do that, we start by defining an abstract semantics of CO₂, which preserves the transitions of a participant $A[P]$, while abstracting those of the context wherein $A[P]$ is run. Building upon the abstract semantics, we then devise an abstract notion of honesty (α -honesty, Definition 11), which neglects the execution context. Theorem 5 states that α -honesty correctly approximates honesty, and that — under certain hypotheses — it is also complete. We then propose a verification technique for α -honesty, and we provide an implementation in Maude. Some experiments have then been carried out; quite notably, our tool has allowed us to determine the dishonesty of a supposedly-honest CO₂ process appeared in [5] (see Example 5).

Because of space limits, we make available online the proofs of all our statements, as well as the Maude implementation, and the experiments made [2].

2 Modelling Contracts

We model contracts as processes in a simple algebra, with internal/external choice and recursion. Compliance between contracts ensures progress, until a successful state is reached. We prove that our model enjoys some relevant properties. First, in each non-final state of a contract there is exactly one participant who is *culpable*, i.e., expected to make the next move (Theorem 1). Furthermore, a participant always recovers from culpability in at most two steps (Theorem 2).

Syntax. We assume a finite set of *participant names* (ranged over by A, B, \dots) and a denumerable set of *atoms* (ranged over by a, b, \dots). We postulate an involution $\text{co}(a)$, also written as \bar{a} , extended to sets of atoms in the natural way. Definition 1 introduces the syntax of contracts. We distinguish between (*unilateral*) contracts c , which model the promised behaviour of a single participant, and *bilateral* contracts γ , which combine the contracts advertised by two participants.

Definition 1. Unilateral contracts are defined by the following grammar:

$$c, d ::= \bigoplus_{i \in \mathcal{I}} a_i ; c_i \mid \sum_{i \in \mathcal{I}} a_i . c_i \mid \text{ready } a.c \mid \text{rec } X. c \mid X$$

where (i) the index set \mathcal{I} is finite; (ii) the “ready” prefix may appear at the top-level, only; (iii) recursion is guarded.

Bilateral contracts γ are terms of the form $A \text{ says } c \mid B \text{ says } d$, where $A \neq B$ and at most one occurrence of “ready” is present. The order of unilateral contracts in γ is immaterial, i.e. $A \text{ says } c \mid B \text{ says } d \equiv B \text{ says } d \mid A \text{ says } c$.

An internal sum $\bigoplus_{i \in \mathcal{I}} a_i ; c_i$ allows to choose one of the branches $a_i ; c_i$, to perform the action a_i , and then to behave according to c_i . Dually, an external sum $\sum_{i \in \mathcal{I}} a_i . c_i$ allows to wait for the other participant to choose one of the branches $a_i . c_i$, then to perform the corresponding a_i and behave according to c_i . Separators $;$ and $.$ allow for distinguishing singleton internal sums $a ; c$ from singleton external sums $a . c$. Empty internal/external sums are denoted with 0 . We will only consider contracts without free occurrences of recursion variables X .

Example 1. An online store A has the following contract: buyers can iteratively add items to the shopping cart (`addToCart`); when at least one item has been added, the client can either `cancel` the order or `pay`; then, the store can accept (`ok`) or decline (`no`) the payment. Such a contract may be expressed as c_A below:

$$\begin{aligned} c_{\text{pay}} &= \text{pay} . (\overline{\text{ok}} ; 0 \oplus \overline{\text{no}} ; 0) \\ c_A &= \text{addToCart} . (\text{rec } Z . \text{addToCart} . Z + c_{\text{pay}} + \text{cancel} . 0) \end{aligned}$$

Instead, a buyer contract could be expressed as:

$$c_B = \text{rec } Z . (\overline{\text{addToCart}} ; Z \oplus \overline{\text{pay}} ; (\text{ok} . 0 + \text{no} . 0))$$

The Maude specification of the syntax of contracts is defined as follows:

```
sorts Atom UniContract Participant AdvContract BiContract
      IGuarded EGuarded IChoice EChoice Var Id RdyContract.
subsort Id < IGuarded < IChoice < UniContract < RdyContract.
subsort Id < EGuarded < EChoice < UniContract < RdyContract.
subsort Var < UniContract.
```

The sorts `IGuarded` and `EGuarded` represent singleton internal/external sums, respectively, while `IChoice` and `EChoice` are for arbitrary internal/external sums. `Id` represents empty sums, and it is a subsort of internal and external sums (either singleton or not). `RdyContract` is for contracts which may have a top-level *ready*, while `AdvContract` is a unilateral contract advertised by some participant.

```
op _- : Atom -> Atom [ctor].
eq -- a:Atom = a:Atom.
op 0 : -> Id [ctor].
op _- : Atom UniContract -> EGuarded [frozen ctor].
op _; : Atom UniContract -> IGuarded [frozen ctor].
op _+ : EChoice EChoice -> EChoice [frozen comm assoc id: 0 ctor].
op _(+)_ : IChoice IChoice -> IChoice [frozen comm assoc id: 0 ctor].
op ready _- : Atom UniContract -> RdyContract [frozen ctor].
op rec _- : Var IChoice -> UniContract [frozen ctor].
op rec _- : Var EChoice -> UniContract [frozen ctor].
op _ says _ : Participant RdyContract -> AdvContract [ctor].
op _ | _ : AdvContract AdvContract -> BiContract [comm ctor].
```

The operator `-` models the involution on atoms, with `eq -- a:Atom = a:Atom`. The other operators are rather standard, and they guarantee that each `UniContract` respects the syntactic constraints imposed by Definition 1.

Semantics. The evolution of bilateral contracts is modelled by $\xrightarrow{\mu}$, the smallest relation closed under the rules in Fig. 1 and under \equiv . The congruence \equiv is the least relation including α -conversion of recursion variables, and satisfying $\text{rec } X. c \equiv c\{\text{rec } X. c/X\}$ and $\bigoplus_{i \in \emptyset} \mathbf{a}_i$; $c_i \equiv \sum_{i \in \emptyset} \mathbf{a}_i . c_i$. The label $\mu = \mathbf{A}$ says \mathbf{a} models \mathbf{A} performing action \mathbf{a} . Hereafter, we shall consider contracts up-to \equiv .

$$\begin{array}{c} \mathbf{A} \text{ says } (\mathbf{a}; c \oplus c') \mid \mathbf{B} \text{ says } (\bar{\mathbf{a}} . d + d') \xrightarrow{\mathbf{A} \text{ says } \mathbf{a}} \mathbf{A} \text{ says } c \mid \mathbf{B} \text{ says } \text{ready } \bar{\mathbf{a}} . d \quad [\text{INTEXT}] \\ \mathbf{A} \text{ says } \text{ready } \mathbf{a} . c \mid \mathbf{B} \text{ says } d \xrightarrow{\mathbf{A} \text{ says } \mathbf{a}} \mathbf{A} \text{ says } c \mid \mathbf{B} \text{ says } d \quad [\text{RDY}] \end{array}$$

Fig. 1. Semantics of contracts (symmetric rules for \mathbf{B} actions omitted)

In rule [INTEXT], participant \mathbf{A} selects the branch \mathbf{a} in an internal sum, and \mathbf{B} is then forced to commit to the corresponding branch $\bar{\mathbf{a}}$ in his external sum. This is done by marking that branch with *ready* $\bar{\mathbf{a}}$, while discarding all the other branches; \mathbf{B} will then perform his action in the subsequent step, by rule [RDY].

In Maude, the semantics of contracts is an almost literal translation of that in Fig. 1 (except that labels are moved to configurations). The one-step transition relation is defined as follows:

$$\begin{array}{l} \text{cr1 [IntExt]: } \mathbf{A} \text{ says } \mathbf{a} ; c (+) c' \mid \mathbf{B} \text{ says } \mathbf{b} . d + d' \\ \quad \Rightarrow \{\mathbf{A} \text{ says } \mathbf{a}\} \mathbf{A} \text{ says } c \mid \mathbf{B} \text{ says } \text{ready } \mathbf{b} . d \quad \text{if } \mathbf{a} = - \mathbf{b}. \\ \text{r1 [Rdy]: } \mathbf{A} \text{ says } \text{ready } \mathbf{a} . c \mid \mathbf{B} \text{ says } d \Rightarrow \{\mathbf{A} \text{ says } \mathbf{a}\} \mathbf{A} \text{ says } c \mid \mathbf{B} \text{ says } d. \end{array}$$

Compliance. Two contracts are *compliant* if, whenever a participant \mathbf{A} wants to choose a branch in an internal sum, then participant \mathbf{B} always offers \mathbf{A} the opportunity to do it. To formalise compliance, we first define a partial function rdy from bilateral contracts to sets of atoms. Intuitively, if the unilateral contracts in γ do not agree on the first step, then $\text{rdy}(\gamma)$ is undefined (i.e. equal to \perp). Otherwise, $\text{rdy}(\gamma)$ contains the atoms which could be fired in the first step.

Definition 2 (Compliance). *Let the partial function rdy be defined as:*

$$\begin{array}{l} \text{rdy} \left(\mathbf{A} \text{ says } \bigoplus_{i \in \mathcal{I}} \mathbf{a}_i ; c_i \mid \mathbf{B} \text{ says } \sum_{j \in \mathcal{J}} \mathbf{b}_j . c_j \right) = \{\mathbf{a}_i\}_{i \in \mathcal{I}} \quad \text{if } \{\mathbf{a}_i\}_{i \in \mathcal{I}} \subseteq \{\bar{\mathbf{b}}_j\}_{j \in \mathcal{J}} \\ \quad \text{and } (\mathcal{I} = \emptyset \implies \mathcal{J} = \emptyset) \\ \text{rdy}(\mathbf{A} \text{ says } \text{ready } \mathbf{a} . c \mid \mathbf{B} \text{ says } d) = \{\mathbf{a}\} \end{array}$$

Then, the compliance relation \bowtie between unilateral contracts is the largest relation such that, whenever $c \bowtie d$:

- (1) $\text{rdy}(A \text{ says } c \mid B \text{ says } d) \neq \perp$
(2) $A \text{ says } c \mid B \text{ says } d \xrightarrow{\mu} A \text{ says } c' \mid B \text{ says } d' \implies c' \bowtie d'$

Example 2. Let $\gamma = A \text{ says } c \mid B \text{ says } d$, where $c = \mathbf{a}; c_1 \oplus \mathbf{b}; c_2$ and $d = \bar{\mathbf{a}}.d_1 + \bar{\mathbf{c}}.d_2$. If the participant A internally chooses to perform \mathbf{a} , then γ will take a transition to $A \text{ says } c_1 \mid B \text{ says } \text{ready } \bar{\mathbf{a}}.d_1$. Suppose instead that A chooses to perform \mathbf{b} , which is not offered by B in his external choice. In this case, $\gamma \not\xrightarrow{A \text{ says } \mathbf{b}}$. We have that $\text{rdy}(\gamma) = \perp$, which does not respect item (1) of Definition 2. Therefore, c and d are *not* compliant.

We say that a contract is *proper* if the prefixes of each summation are pairwise distinct. The next lemma states that each proper contract has a compliant one.

Lemma 1. *For all proper contracts c , there exists d such that $c \bowtie d$.*

Definition 2 cannot be directly exploited as an algorithm for checking compliance. Lemma 2 gives an alternative, model-checkable characterisation of \bowtie .

Lemma 2. *For all bilateral contracts $\gamma = A \text{ says } c \mid B \text{ says } d$:*

$$c \bowtie d \iff (\forall \gamma'. \gamma \twoheadrightarrow^* \gamma' \implies \text{rdy}(\gamma') \neq \perp)$$

In Maude, the compliance relation is defined as suggested by Lemma 2. The predicate `isBottom` is true for a contract γ whenever $\text{rdy}(\gamma) = \perp$. The operator `<>` used below allows for the transitive closure of the transition relation. The relation `c |X| d` is implemented by verifying that the contract $A \text{ says } c \mid B \text{ says } d$ satisfies the LTL formula $\Box \neg \text{isBottom}$. This is done through the Maude model checker.

```
eq <{1} g> |= isBottom = is rdy(g) eq bottom.
op _|X|_ : UniContract UniContract -> Bool.
eq c |X| d = modelCheck(<A says c | B says d>, [] ~isBottom) == true.
```

Example 3. Recall the store contract c_A in Example 1. Its Maude version is:

```
op Z : -> Var.
ops addToCart pay ok no cancel : -> Atom.
ops CA CPay CB : -> UniContract.
eq CPay = pay . (- ok ; 0 (+) - no ; 0).
eq CA = addToCart . (rec Z . addToCart . Z + CPay + cancel . 0).
```

Instead, the Maude implementation of the buyer contract c_B in Example 1 is:

```
eq CB = rec Z . ( - addToCart ; Z (+) - pay ; (ok . 0 + no . 0) ).
```

We can verify with Maude that CA and CB are *not* compliant:

```
red CA |X| CB.
result Bool: false
```

The problem is that CB may choose to pay even when the cart is empty. We can easily fix the buyer contract as follows, and then obtain compliance:

```
red CA |X| (- addToCart ; CB).
result Bool: true
```

Culpability. We now tackle the problem of determining who is expected to make the next step for the fulfilment of a bilateral contract. We call a participant A *culpable* in γ if she is expected to perform some actions so to make γ progress.

Definition 3. *A participant A is culpable in γ ($A \dot{\curvearrowright} \gamma$ in symbols) iff $\gamma \xrightarrow{A \text{ says } a}$ for some a . When A is not culpable in γ we write $A \dot{\curvearrowleft} \gamma$.*

Theorem 1 below establishes that, when starting with compliant contracts, exactly one participant is culpable in a bilateral contract. The only exception is $A \text{ says } 0 \mid B \text{ says } 0$, which represents a successfully terminated interaction, where nobody is culpable.

Theorem 1. *Let $\gamma = A \text{ says } c \mid B \text{ says } d$, with $c \bowtie d$. If $\gamma \rightarrow^* \gamma'$, then either $\gamma' = A \text{ says } 0 \mid B \text{ says } 0$, or there exists a unique culpable in γ' .*

The following theorem states that a participant is always able to recover from culpability by performing some of her duties. This requires at most two steps.

Theorem 2 (Contractual exculpation). *Let $\gamma = A \text{ says } c \mid B \text{ says } d$. For all γ' such that $\gamma \rightarrow^* \gamma'$, we have that:*

- (1) $\gamma' \not\curvearrowright \implies A \dot{\curvearrowleft} \gamma' \text{ and } B \dot{\curvearrowleft} \gamma'$
- (2) $A \dot{\curvearrowright} \gamma' \implies \forall \gamma'' . \gamma' \rightarrow \gamma'' \implies \begin{cases} A \dot{\curvearrowright} \gamma'' , \text{ or} \\ \forall \gamma''' . \gamma'' \rightarrow \gamma''' \implies A \dot{\curvearrowright} \gamma''' \end{cases}$

Item (1) of Theorem 2 says that, in a stuck contract, no participant is culpable. Item (2) says that if A is culpable, then she can always exculpate herself in *at most* two steps, i.e.: one step if A has an internal choice, or a *ready* followed by an external choice; two steps if A has a *ready* followed by an internal choice.

We specify culpability in Maude as follows. The formula $\{1\} g \mid = \text{--}A\text{--}\gg$ is true whenever g has been reached by some transitions of A . The participant A is culpable in g , written $A :C g$, if g satisfies the LTL formula $0 \text{--}A\text{--}\gg$ (where 0 is the “next” operator of LTL). This is verified through the Maude model checker.

```

op --_-->> : Participant -> Prop.
eq {A says a} g \mid = -- A -->> = true.
eq {1} g \mid = -- A -->> = false [owise].
op _ :C _ : Participant BiContract -> Bool.
eq A :C g = modelCheck(g, 0 -- A -->>) == true.
    
```

3 Modelling Contracting Processes

We model agents and systems through the process calculus CO_2 [3], which we instantiate with the contracts introduced in Sect. 2. The primitives of CO_2 allow agents to advertise contracts, to open sessions between agents with compliant contracts, to execute them by performing some actions, and to query contracts.

Syntax. Let \mathcal{V} and \mathcal{N} be disjoint sets of *session variables* (ranged over by x, y, \dots) and *session names* (ranged over by s, t, \dots). Let u, v, \dots range over $\mathcal{V} \cup \mathcal{N}$, and \mathbf{u}, \mathbf{v} range over $2^{\mathcal{V} \cup \mathcal{N}}$.

Definition 4. *The syntax of CO_2 is given as follows:*

$$\begin{array}{l} \text{Systems } S ::= \mathbf{0} \quad \left| \quad \mathbf{A}[P] \quad \left| \quad s[\gamma] \quad \left| \quad S \mid S \quad \left| \quad (u)S \quad \left| \quad \{\downarrow_x c\}_{\mathbf{A}} \right. \right. \right. \\ \text{Processes } P ::= \sum_i \pi_i.P_i \quad \left| \quad P \mid P \quad \left| \quad (u)P \quad \left| \quad X(\mathbf{u}) \right. \right. \right. \\ \text{Prefixes } \pi ::= \tau \quad \left| \quad \text{tell } \downarrow_u c \quad \left| \quad \text{do}_u a \quad \left| \quad \text{ask}_u \phi \right. \right. \right. \end{array}$$

Systems are the parallel composition of *participants* $\mathbf{A}[P]$, *delimited systems* $(u)S$, *sessions* $s[\gamma]$ and *latent contracts* $\{\downarrow_x c\}_{\mathbf{A}}$. A latent contract $\{\downarrow_x c\}_{\mathbf{A}}$ represents a contract c (advertised by \mathbf{A}) which has not been stipulated yet; upon stipulation, the variable x will be instantiated to a fresh session name. We assume that, in a system of the form $(\mathbf{u})(\mathbf{A}[P] \mid \mathbf{B}[Q]) \mid \dots$, $\mathbf{A} \neq \mathbf{B}$. We denote with \mathbf{K} a special participant name (playing the role of contract broker) such that, in each system $(\mathbf{u})(\mathbf{A}[P] \mid \dots)$, $\mathbf{A} \neq \mathbf{K}$. We allow for prefix-guarded finite sums of processes, and write $\pi_1.P_1 + \pi_2.P_2$ for $\sum_{i \in \{1,2\}} \pi_i.P_i$, and $\mathbf{0}$ for $\sum_{\emptyset} P$. Recursion is allowed only for processes; we stipulate that each process identifier X has a unique defining equation $X(x_1, \dots, x_j) \stackrel{\text{def}}{=} P$ such that $\text{fv}(P) \subseteq \{x_1, \dots, x_j\} \subseteq \mathcal{V}$, and each occurrence of process identifiers in P is prefix-guarded. We will sometimes omit the arguments of $X(\mathbf{u})$ when they are clear from the context.

Prefixes include silent action τ , contract advertisement $\text{tell } \downarrow_u c$, action execution $\text{do}_u a$, and contract query $\text{ask}_u \phi$ (where ϕ is an LTL formula on γ). In each prefix $\pi \neq \tau$, u refers to the target session involved in the execution of π .

In Maude, we translate the syntax of CO_2 almost literally. Here we just show the sorts used; see [2] for the full details.

```
sorts System Process Prefix SessionName SessionVariable SessionIde
      GuardProc Sum IdeVec ProcIde ParamList.
subsort SessionName < SessionIde < IdeVec.
subsort Qid < SessionVariable < SessionIde < IdeVec.
subsort GuardProc < Sum < Process.
subsort SessionIde < ParamList.
```

The sort `SessionIde` is a super sort of both `SessionVariable` and `SessionName`. Session variables can be of sort `Qid`; session names can not. Sort `IdeVec` models sets of `SessionIde` (used as syntactic sugar for delimitations), while `ParamList` models vectors of `SessionIde` (used for parameters of defining equations).

Semantics. The CO_2 semantics is formalised by the relation $\xrightarrow{\mu}$ in Fig. 3, where $\mu \in \{\mathbf{A}:\pi \mid \mathbf{A} \neq \mathbf{K}\} \cup \{\mathbf{K}:\text{fuse}\}$. We will consider processes and systems up-to the congruence relation \equiv in Fig. 2. The axioms for \equiv are fairly standard — except the last one: it collects garbage terms possibly arising from variable substitutions.

Rule [TAU] just fires a τ prefix. Rule [TELL] advertises a latent contract $\{\downarrow_x c\}_{\mathbf{A}}$. Rule [FUSE] finds *agreements* among the latent contracts: it happens when there exist $\{\downarrow_x c\}_{\mathbf{A}}$ and $\{\downarrow_y d\}_{\mathbf{B}}$ such that $\mathbf{A} \neq \mathbf{B}$ and $c \bowtie d$. Once the agreement is reached, a fresh session containing $\gamma = \mathbf{A} \text{ says } c \mid \mathbf{B} \text{ says } d$ is created. Rule [DO] allows a participant \mathbf{A} to perform an action in the session s containing γ (which,

commutative monoidal laws for $|$ on processes and systems

$$\begin{aligned} A[(v)P] &\equiv (v)A[P] & Z | (u)Z' &\equiv (u)(Z | Z') \text{ if } u \notin \text{fv}(Z) \cup \text{fn}(Z) \\ (u)(v)Z &\equiv (v)(u)Z & (u)Z &\equiv Z \text{ if } u \notin \text{fv}(Z) \cup \text{fn}(Z) & \{\downarrow_s c\}_A &\equiv \mathbf{0} \end{aligned}$$

Fig. 2. Structural equivalence for CO_2 (Z, Z' range over systems or processes).

$$\begin{aligned} & \frac{}{A[\tau.P + P' | Q] \xrightarrow{A: \tau} A[P | Q]} \quad \text{[TAU]} \\ & \frac{}{A[\text{tell } \downarrow_u c.P + P' | Q] \xrightarrow{A: \text{tell } \downarrow_u c} A[P | Q] | \{\downarrow_u c\}_A} \quad \text{[TELL]} \\ & \frac{c \bowtie d \quad \gamma = A \text{ says } c | B \text{ says } d \quad \sigma = \{s/x, y\} \quad s \text{ fresh}}{(x, y)(S | \{\downarrow_x c\}_A | \{\downarrow_y d\}_B) \xrightarrow{K: \text{fuse}} (s)(S\sigma | s[\gamma])} \quad \text{[FUSE]} \\ & \frac{\gamma \xrightarrow{A \text{ says } a} \gamma'}{A[\text{do}_s a.P + P' | Q] | s[\gamma] \xrightarrow{A: \text{do}_s a} A[P | Q] | s[\gamma']} \quad \text{[DO]} \\ & \frac{\gamma \vdash \phi}{A[\text{ask}_s \phi.P + P' | Q] | s[\gamma] \xrightarrow{A: \text{ask}_s \phi} A[P | Q] | s[\gamma]} \quad \text{[ASK]} \\ & \frac{X(\mathbf{u}) \stackrel{\text{def}}{=} P \quad A[P\{v/u\} | Q] | S \xrightarrow{\mu} S'}{A[X(v) | Q] | S \xrightarrow{\mu} S'} \quad \text{[DEF]} \qquad \frac{S \xrightarrow{\mu} S'}{S | S'' \xrightarrow{\mu} S' | S''} \quad \text{[PAR]} \\ & \frac{S \xrightarrow{A: \pi} S'}{(u)S \xrightarrow{A: \text{del}_u(\pi)} (u)S'} \quad \text{[DEL]} \quad \text{where } \text{del}_u(\pi) = \begin{cases} \tau & \text{if } u \in \text{fnv}(\pi) \\ \pi & \text{otherwise} \end{cases} \end{aligned}$$

Fig. 3. Reduction semantics of CO_2 .

accordingly, evolves to γ'). Rule [ASK] allows **A** to proceed only if the contract γ at session s satisfies the property ϕ . The last three rules are mostly standard. In rule [DEL] the label π fired in the premise becomes τ in the consequence, when π contains the delimited name/variable. This transformation is defined by the function $\text{del}_u(\pi)$, where the set $\text{fnv}(\pi)$ contains the free names/variables in π . For instance, $(x)A[\text{tell } \downarrow_x c.P] \xrightarrow{A: \tau} (x)(A[P] | \{\downarrow_x c\}_A)$. Here, it would make little sense to have the label $A: \text{tell } \downarrow_x c$, as x (being delimited) may be α -converted.

Implementing in Maude the semantics of CO_2 is almost straightforward [19]; here we show only the main rules (see [2] for the others). Rule [DO] uses the transition relation \Rightarrow on bilateral contracts. Rule [ASK] exploits the Maude model checker to verify if the bilateral contract \mathbf{g} satisfies the LTL formula \mathbf{phi} . Rule [FUSE] uses the operator $|X|$ to check compliance between the contracts c and d , then creates the session $s[A \text{ says } c | B \text{ says } d]$ (with s fresh), and

finally applies the substitution $\{s / x\}\{s / y\}$ (delimitations are dealt with as in Fig. 3).

```

crl [Do] : A[do s a . P + P' | Q] | s[g] => {A : do s a} (A[P | Q] | s[g'])
          if g => {A says a} g'.

crl [Ask] : A[ask s phi . P + P' | Q] | s[g] => {A : ask s phi} A[P | Q]
          if g |- phi.

crl [Fuse] : (uVec , vVec) ({x c}A | {y d}B | S) => {K : fuse}
          (s , vVec) (s[A says c | B says d] | S{s / x}{s / y})
          if uVec == (x , y) / c |X| d / s := fresh(0 , S).

```

4 Honesty

A remarkable feature of CO₂ is that it allows for writing *dishonest* agents which do not keep their promises. Intuitively, a participant is honest if she always fulfils her contractual obligations, in all possible contexts. Below we formalise the notion of honesty, by slightly adapting the one appeared in [3]. Then, we show how we verify in Maude a weaker notion, i.e. honesty *in a given context*.

We start by defining the set $O_s^A(S)$ of *obligations* of A at s in S . Whenever A is culpable at some session s , she has to fire one of the actions in $O_s^A(S)$.

Definition 5. We define the set of atoms $O_s^A(S)$ as:

$$O_s^A(S) = \left\{ a \mid \exists \gamma, S' . S \equiv s[\gamma] \mid S' \text{ and } \gamma \xrightarrow{A \text{ says } a} \right\}$$

We say that A is culpable at s in S iff $O_s^A(S) \neq \emptyset$.

The set of atoms $RD_s^A(S)$ (“Ready Do”) defined below comprises all the actions that A can perform at s in one computation step within S (note that, by rule [DEL], if s is a bound name then $RD_s^A(S) = \emptyset$). The set $WRD_s^A(S)$ (“Weak Ready Do”) contains all the actions that A may possibly perform at s after a finite sequence of transitions of A not involving any do at s .

Definition 6. For all S , A and s , we define the sets of atoms:

$$RD_s^A(S) = \left\{ a \mid \exists S' . S \xrightarrow{A: do_s a} S' \right\}$$

$$WRD_s^A(S) = \left\{ a \mid \exists S' . S \xrightarrow{A: \neq do_s}^* S' \wedge a \in RD_s^A(S') \right\}$$

where we write $S \xrightarrow{A: \neq do_s} S'$ if $\exists \pi . S \xrightarrow{A: \pi} S' \wedge \forall a . \pi \neq do_s a$.

A participant is *ready* if she can fulfil some of her obligations. To check if A is ready in S , we consider all the sessions s in S involving A. For each of them, we check that some obligations of A at s are exposed after some steps of A *not* preceded by other do_s of A. A[P] is honest *in a given system* S when A is ready in all evolutions of A[P] | S . Then, A[P] is honest when she is honest in *all* S .

Definition 7 (Honesty). *We say that:*

1. S is A -free iff it has no latent/stipulated contracts of A , nor processes of A
2. A is ready in S iff $S \equiv (\mathbf{u})S' \wedge O_s^A(S') \neq \emptyset \implies \text{WRD}_s^A(S') \cap O_s^A(S') \neq \emptyset$
3. P is honest in S iff $\forall A : (S \text{ is } A\text{-free} \wedge A[P] \mid S \rightarrow^* S') \implies A \text{ is ready in } S'$
4. P is honest iff, for all S , P is honest in S

We have implemented items 2 and 3 of the above definition in Maude (item 4 is dealt with in the next section). CO_2 can simulate Turing machines [5], hence reachability in CO_2 is undecidable, and consequently WRD, readiness and honesty are undecidable as well. To recover decidability, we then restrict to finite state processes: roughly, these are the processes with neither delimitations nor parallel compositions under process definitions.

In Maude we verify readiness in a session s by searching if A can reach (with her moves only), a state which allows for a do_s a move, for some a .

```
op ready? : Participant SessionName System Module -> Bool.
eq ready?(A,s,S,M:Module) = metaSearch(M:Module, upTerm(< S > A s),
  '<_>_['S1:System, upTerm(A), upTerm(s)],
  'S1:System => '[_]'_['1:SLabel,'S2:System] /\
  '[_]'_['upTerm(A),'do_['upTerm(s),'a:Atom]] := '1:SLabel,
  '*, unbounded, 0 ) =/= failure.
```

We start the search from the term $\langle S \rangle A s$, whose meta-representation is obtained through the `upTerm` function. The search is performed according to the A -solo semantics of CO_2 (see Definition 10), which blocks all do at s . This is done by the operator `<_>_`. Then, we look for reachable systems $S1$ where A can fire a do at s . If the search succeeds, `ready?` returns true. Note that if A has no obligations at s in S , `ready?` returns false — uncoherently with Definition 7. To correctly check readiness, we define the function `ready` (see [2]), which invokes `ready?` only when $O_s^A(S) \neq \emptyset$.

Verifying honesty in a context is done similarly. We use `metaSearch` to check that A is ready in all reachable states. The operator `<_>` gives the CO_2 semantics.

```
op search-honest-ctx : Participant System Module -> ResultTriple?.
eq search-honest-ctx(A,S,M:Module) = metaSearch(M:Module, upTerm(< S >),
  '<_>['S:System], 'ready[upTerm(A), 'S:System,'S:System, upTerm(M:Module)]
  = 'false.Bool, '*, unbounded, 0).
op honest-ctx : Participant System Module -> Result.
ceq honest-ctx (A , S , M:Module) = true
  if search-honest-ctx (A , S , M:Module) == failure.
ceq honest-ctx (A , S , M:Module) = downTerm (T:Term , < (0).System > )
  if {T:Term,Ty:Type,S:Substitution} := search-honest-ctx (A,S,M:Module).
```

Example 4. A travel agency A queries in parallel an airline ticket broker F and a hotel reservation service H in order to organise a trip for some user U . The agency first requires U to pay, and then chooses either to commit the reservation or to issue a refund (contract CU). When querying the ticket broker (contract CF), the agency first receives a quotation, and then chooses either to commit and pay the ticket, or to abort the transaction. The contract CH between A and H is similar.

```

eq CU = pay . (commit ; 0 (+) refund ; 0).
eq CF = ticket . ( commitF ; payF ; 0 (+) abortF ; 0).
eq CH = hotel . ( commitH ; payH ; 0 (+) abortH ; 0).

```

In addition to the contracts above, the agency should respect the following constraints: (a) the agency refunds U only if both the transactions with F and H are aborted; (b) A pays the ticket and the hotel reservation only after it has committed the transaction with U ; (c) either both the transactions with F or H are committed, or they are both aborted. A possible specification in Maude respecting the above constraints is given by the following process P :

```

eq P = ( xu , xf , xh ) ( tell xu CU . do xu pay .
  ( (tell xf CF . PF) | (tell xh CH . PH) | PU ) ).

eq PF = do xf ticket . (do xh commitH . 0 + do xf abortF . 0).
eq PH = do xh hotel . (do xf commitF . 0 + do xh abortH . 0).

eq PU = ask xh ([ ] ~ payH) . do xu refund . 0 +
  t . do xu commit . (do xf payF . 0 | do xh payH . 0).

```

The process P first opens a session with U , and then advertises the contracts CF and CH , and in parallel executes PU . The process PF gets the ticket quotation, then either commits the hotel reservation, or aborts the flight reservation. Dually, PH gets the hotel quotation, then either commits the flight reservation, or aborts the hotel reservation. Note that the two choices in PF and PH ensure that constraint (c) above is satisfied: e.g., if PF fires the `commitH` (resp. `abortF`) prefix, the `abortH` (resp. `commitF`) branch in PH is disabled, and only `commitF` (resp. `abortH`) can be selected. The process PU checks if a refund is due to U . When the atom `payH` is no longer reachable in session xh , the `ask` passes, and the refund is issued. This guarantees constraint (a). In the τ -branch, PU commits the transaction with U , and then proceeds to pay both F and H . This satisfies constraint (b). Note that it may happen that PU chooses to `commit` even when CF or CH are not stipulated. Although this behaviour is conceptually wrong, it does not affect honesty. Indeed, honesty does not consider the domain-specific constraints among actions (e.g. (a), (b), (c) above), but only that the advertised contracts are respected.

We have experimented the function `honest-ctx` by inserting P in some contexts S where all the other participants U , F and H are honest (see [2] for details). The Maude model checker has correctly determined that P is honest in S .

```

red honest-ctx(A , S , ['TRAVEL-AGENCY-CTX]).
rewrites: 53950741 in 38062ms cpu (38058ms real) (1417429 rewrites/second)
result Bool: true

```

Even though we conjecture that P is honest (in all contexts), we anticipate here that the verification technique proposed in Sect. 5 does not classify P as honest. This is because the analysis is (correct but) not complete in the presence of `ask`: indeed, the precise behaviour of an `ask` is lost by the analysis, because it abstracts from the contracts of the context.

5 Model Checking Honesty

We now address the problem of automatically verifying honesty. As mentioned in Sect. 1, this is a desirable goal, because it alerts system designers before they deploy services which could violate contracts at run-time (so possibly incurring in sanctions). Since honesty is undecidable in general [5], our goal is a verification technique which safely over-approximates honesty, i.e. it never classifies a process as honest when it is not. The first issue is that Definition 7 requires readiness to be preserved in all possible contexts, and there is an *infinite* number of such contexts. To overcome this problem, we present below an *abstract* semantics of CO₂ which preserves the honesty property, while neglecting the actual context where the process $A[P]$ is executed.

The definition of the abstract semantics of CO₂ is obtained in two steps. First, we provide the projections from concrete contracts/systems to the abstract ones. Then, we define the semantics of abstract contracts and systems, and we relate the abstract semantics with the concrete one. The abstraction is always parameterised in the participant A the honesty of which is under consideration.

The abstraction $\alpha_A(\gamma)$ of a bilateral contract $\gamma = A \text{ says } c \mid B \text{ says } d$ (Definition 8 below) is either c , or $ctx.c$ when d has a *ready*.

Definition 8. For all γ , we define the abstract contract $\alpha_A(\gamma)$ as:

$$\alpha_A(A \text{ says } c \mid B \text{ says } d) = \begin{cases} c & \text{if } d \text{ is ready-free} \\ ctx \ a.c & \text{if } d = \text{ready } a.d \end{cases}$$

We now define the abstraction α_A of concrete systems, which just discards all the components not involving A , and projects the contracts involving A .

Definition 9. For all A, S we define the abstract system $\alpha_A(S)$ as:

$$\begin{aligned} \alpha_A(A[P]) &= A[P] & \alpha_A(s[\gamma]) &= s[\alpha_A(\gamma)] & \alpha_A(\{\downarrow_x c\}_A) &= \{\downarrow_x c\}_A \\ \alpha_A(S \mid S') &= \alpha_A(S) \mid \alpha_A(S') & \alpha_A((u)S) &= (u)(\alpha_A(S)) & \alpha_A(S) &= \mathbf{0}, \text{ otherwise} \end{aligned}$$

Abstract semantics. For all participants A , the abstract LTSs $\xrightarrow{\ell}_A$ and $\xrightarrow{\mu}_A$ on abstract contracts and systems, respectively, are defined by the rules in Fig. 4. Labels ℓ are atoms, with or without the special prefix *ctx* — which indicates a contractual action performed by the context. Labels μ are either *ctx* or they have the form $A: \pi$, where A is the participant in \rightarrow_A , and π is a CO₂ prefix.

Rules for abstract contracts (first row in Fig. 4) are simple: in an internal sum, A chooses a branch; in an external sum, the choice is made by the context; in a *ready a.c* the atom a is fired. The rightmost rule handles a *ready* in the context contract. For abstract systems, some rules are similar to the concrete ones, hence we discuss only the most relevant ones. Rule $[\alpha\text{-Do}]$ involves the abstract transitions of contracts. The behaviour of abstract systems also considers context actions, labelled with *ctx*. If $c \vdash \phi$, then the **ask** ϕ passes, independently from the context (rule $[\alpha\text{-Ask}]$). If $c \not\vdash \neg\phi$, then the **ask** ϕ may pass or not, depending and

the context (rule $[\alpha\text{-ASKCTX}]$). Rule $[\alpha\text{-FUSE}]$ says that a latent contract of \mathbf{A} may always be fused (the context may choose whether this is the case or not). The context may also decide whether to perform actions within sessions ($[\alpha\text{-DoCTX}]$). Unobservable context actions are modelled by rules $[\alpha\text{-CTX}]$ and $[\alpha\text{-DELCTX}]$.

$$\begin{array}{c}
\mathbf{a}; c \oplus c' \xrightarrow{\mathbf{a}}_{\mathbf{A}} \text{ctx } \bar{\mathbf{a}}.c \quad \mathbf{a}.c + c' \xrightarrow{\text{ctx}:\bar{\mathbf{a}}}_{\mathbf{A}} \text{ready } \mathbf{a}.c \quad \text{ready } \mathbf{a}.c \xrightarrow{\mathbf{a}}_{\mathbf{A}} c \quad \text{ctx } \mathbf{a}.c \xrightarrow{\text{ctx}:\mathbf{a}}_{\mathbf{A}} c \\
\frac{c \xrightarrow{\mathbf{a}}_{\mathbf{A}} c'}{\mathbf{A}[\text{do}_s \mathbf{a}.P + P' \mid Q] \mid s[c] \xrightarrow{\mathbf{A}:\text{do}_s \mathbf{a}}_{\mathbf{A}} \mathbf{A}[P \mid Q] \mid s[c']} \quad [\alpha\text{-Do}] \\
\frac{s \text{ fresh}}{(x)(\tilde{S} \mid \{\downarrow_x c\}_{\mathbf{A}}) \xrightarrow{\text{ctx}}_{\mathbf{A}} (s)(s[c] \mid \tilde{S}\{s/x\})} \quad [\alpha\text{-FUSE}] \\
\frac{c \vdash \phi}{\mathbf{A}[\text{ask}_s \phi.P + P' \mid Q] \mid s[c] \xrightarrow{\mathbf{A}:\text{ask}_s \phi}_{\mathbf{A}} \mathbf{A}[P \mid Q] \mid s[c]} \quad [\alpha\text{-ASK}] \\
\frac{c \not\vdash \neg\phi}{\mathbf{A}[\text{ask}_s \phi.P + P' \mid Q] \mid s[c] \xrightarrow{\text{ctx}}_{\mathbf{A}} \mathbf{A}[P \mid Q] \mid s[c]} \quad [\alpha\text{-ASKCTX}] \\
\frac{c \xrightarrow{\text{ctx}}_{\mathbf{A}} c'}{s[c] \xrightarrow{\text{ctx}}_{\mathbf{A}} s[c']} \quad [\alpha\text{-DoCTX}] \quad S \xrightarrow{\text{ctx}}_{\mathbf{A}} S \quad [\alpha\text{-CTX}] \quad \frac{\tilde{S} \xrightarrow{\text{ctx}}_{\mathbf{A}} \tilde{S}'}{(u)\tilde{S} \xrightarrow{\text{ctx}}_{\mathbf{A}} (u)\tilde{S}'} \quad [\alpha\text{-DELCTX}]
\end{array}$$

Fig. 4. Abstract LTSs for contracts and systems (full set of rules in [2]).

To check if $\mathbf{A}[P]$ is honest, we must only consider those \mathbf{A} -free contexts not already containing advertised/stipulated contracts of \mathbf{A} . Such systems will always evolve to a system which can be split in two parts: an \mathbf{A} -solo system $S_{\mathbf{A}}$ containing the process of \mathbf{A} , the contracts advertised by \mathbf{A} and all the sessions containing contracts of \mathbf{A} , and an \mathbf{A} -free system S_{ctx} .

Definition 10. *We say that a system S is \mathbf{A} -solo iff one of the following holds:*

$$\begin{array}{llll}
S \equiv \mathbf{0} & S \equiv \mathbf{A}[P] & S \equiv s[\mathbf{A} \text{ says } c \mid \mathbf{B} \text{ says } d] & S \equiv \{\downarrow_x c\}_{\mathbf{A}} \\
S \equiv S' \mid S'' & \text{where } S' \text{ and } S'' \text{ } \mathbf{A}\text{-solo} & S \equiv (u)S' & \text{where } S' \text{ } \mathbf{A}\text{-solo}
\end{array}$$

We say that S is \mathbf{A} -safe iff $S \equiv (s)(S_{\mathbf{A}} \mid S_{\text{ctx}})$, with $S_{\mathbf{A}}$ \mathbf{A} -solo and S_{ctx} \mathbf{A} -free.

The following theorems establish the relations between the concrete and the abstract semantics of CO_2 . Theorem 3 states that the abstraction is *correct*, i.e. for each concrete computation there exists a corresponding abstract computation. Theorem 4 states that the abstraction is also *complete*, provided that a process has neither *ask* nor non-proper contracts.

Theorem 3. *For all A-safe systems S , and for all concrete traces η :*

$$S \xrightarrow{\eta}^* S' \implies \exists \tilde{\eta} : \alpha_A(S) \xrightarrow{\tilde{\eta}}^* \alpha_A(S')$$

Furthermore, if η is A-solo and S is ask-free, then $\eta = \tilde{\eta}$.

Theorem 4. *For all ask-free abstract system \tilde{S} with proper contracts only:*

$$\tilde{S} \rightarrow_A^* \tilde{S}' \implies \exists S, S' \text{ A-safe. } \alpha_A(S) = \tilde{S} \wedge S \rightarrow^* S' \wedge \alpha_A(S') = \tilde{S}'$$

The abstract counterparts of Ready Do, Weak Ready Do, and readiness are defined as expected, by using the abstract semantics instead of the concrete one (see [2] for details). The notion of honesty for abstract systems, namely α -honesty, follows the lines of that of honesty in Definition 7.

Definition 11 (α -honesty). *We say that P is α -honest iff for all \tilde{S} such that $A[P] \rightarrow_A^* \tilde{S}$, A is ready in \tilde{S} .*

The main result of this paper follows. It states that α -honesty is a sound approximation of honesty, and — under certain conditions — it is also complete.

Theorem 5. *If P is α -honest, then P is honest. Conversely, if P is honest, ask-free, and has proper contracts only, then P is α -honest.*

In Maude, we implement abstract semantics for system and contracts for one-step transitions. We obtain their transitive closure, discarding labels, with the operator `<_>`. The function `ready` in `search-honest` computes abstract readiness.

```

op search-honest : Process Module -> ResultTriple?.
eq search-honest (P , M:Module) = metaSearch(M:Module , upTerm(< A[P] > ) ,
  '<_>['S:System] , 'ready['S:System , 'S:System , upTerm(M:Module)])
  = 'false.Bool , '* , unbounded , 0).

op honest : Process Module -> Result.
ceq honest (P , M:Module) = true if search-honest (P,M:Module) == failure . ceq honest (P ,
M:Module) = downTerm (T:Term , < (0).System > )
  if {T:Term , Ty:Type , S:Substitution} := search-honest (P , M:Module).

```

Honesty is checked by searching for states such that A is *not* ready. If the search fails, then A is honest. As in Sect. 4, this function is decidable for finite state processes, i.e. those without delimitation/parallel under process definitions. The following example shows a process which was erroneously classified as honest in [5]. The Maude model checker has determined the dishonesty of that process, and by exploiting the Maude tracing facilities we managed to fix it.

Example 5. A store A offers buyers two options: `clickPay` or `clickVoucher`. If a buyer B chooses `clickPay`, A requires a payment (`pay`) otherwise A checks the validity of the voucher with V , an online voucher distribution system. If V validates the voucher (`ok`), B can use it (`voucher`), otherwise (`no`) B must `pay`. We specify in Maude the contracts `CB` (between A and B) and `CV` (between A and V) as:

```

eq CB = clickPay . pay . 0 +
        clickVoucher . (- reject ; pay . 0 (+) - accept ; voucher . 0).
eq CV = ok . 0 + no . 0.

```

We can specify in Maude a CO₂ process for A as follows:

```

eq P = (x)(tell x CB . (do x clickPay . do x pay . 0 +
                        do x clickVoucher . ((y) tell y CV . Q))).
eq Q = do y ok . do x - accept . do x voucher . 0 +
        do y no . do x - reject . do x pay . 0 + R.
eq R = t . (do x - reject . do x pay . 0).

```

Variables x and y in P correspond to two separate sessions, where A respectively interacts with B and V . The advertisement of CV causally depends on the stipulation of the contract CB , because A must fire `clickVoucher` before `tell y CV`. In process Q the store waits for the answer of V : if V validates the voucher (first branch), then A accepts it from B ; otherwise (second branch), A requires B to pay. The third branch R allows A to fire a τ action, and then reject the voucher. The intuition is that τ models a timeout, to deal with the fact that CV might not be stipulated. When we check the honesty of P with Maude, we obtain:

```

red honest(P , ['STORE-VOUCHER]).
rewrites: 31649 in 72ms cpu (77ms real) (439545 rewrites/second)
result TSystem: < ($ 0,$ 1)(A[do $ 0 - reject . do $ 0 pay . (0).Sum] |
  $ 0[- accept ; voucher . 0(+)- reject ; pay . 0] | $ 1[ready ok . 0]) >

```

This means that the process P is dishonest: actually, the output provides a state where A is not ready. There, A must do `ok` in session y ($\$1$), while A is only ready to do a `-reject` at session x ($\$0$). This problem occurs when the branch R is chosen. To recover honesty, it suffices to replace R with the following process R' :

```

eq R' = t . (do x - reject . do x pay . 0 | (do y no . 0 + do y ok . 0)).
red honest(P' , ['STORE-VOUCHER]).
rewrites: 44009 in 32ms cpu (30ms real) (1375195 rewrites/second)
result Bool: true

```

6 Conclusions

We have described an executable specification in Maude of a calculus for contract-oriented systems. This has been done in two steps. First, we have specified a model for contracts, and we have formalised in Maude their semantics, and the crucial notions of compliance and culpability (Sect. 2). This specification has been exploited in Sect. 3 to implement in Maude the calculus CO₂ [4]. Then, we have considered the problem of honesty [5], i.e. that of deciding when a participant always respects the contracts she advertises, in all possible contexts (Sect. 4). Writing honest processes is not a trivial task, especially when multiple sessions are needed for realising a contract (see e.g. Example 4 and Example 5). We have then devised a sound verification technique for deciding when a participant is honest, and we have provided an implementation of this technique in Maude (Sect. 5).

Related work. Rewriting logic [12] has been successfully used for more than two decades as a semantic framework wherein many different programming models and logics are naturally formalised, executed and analysed. Just by restricting to models for concurrency, there exist Maude specifications and tools for CCS [17], the π -calculus [16], Petri nets [15], Erlang [14], Klaim [18], adaptive systems [7], *etc.* A more comprehensive list of calculi, programming languages, tools and applications implemented in Maude is collected in [13].

The contract model presented in Sect. 2 is a refined version of the one in [5], which in turn is an alternative formalisation of the one in [8]. Our version is simpler and closer to the notion of *session behaviour* [1], and enjoys several desirable properties. Theorem 1 establishes that only one participant may be culpable in a bilateral contract, whereas in [5] both participants may be culpable, e.g. in $A \text{ says } a; c \mid B \text{ says } \bar{a}; d$. In our model, if both participants have an internal (or external) choice, then their contracts are *not* compliant, whereas e.g. $a.c$ and $\bar{a}.d$ (both external choices) are compliant in [5, 8] whenever c and d are compliant. The exculpation property established by Theorem 2 is stronger than the corresponding one in [5]. There, a participant A is guaranteed to exculpate herself by performing (at most) two consecutive actions of A , while in our model two any actions (of *whatever* participant) suffice.

As far as we know, the concept of *contract-oriented computing* (in the meaning used in this paper) has been introduced in [6]. CO₂, a contract-agnostic calculus for contract-oriented computing, has been instantiated with several contract models — both bilateral [3, 5] and multiparty [4, 11]. Here we have instantiated it with the contracts in Sect. 2. A minor difference w.r.t. [3, 5, 11] is that here we no longer have *fuse* as a language primitive, but rather the creation of fresh sessions is performed non-deterministically by the context (rule [FUSE]). This is equivalent to assume a contract broker which collects all contracts, and may establish sessions when compliant contracts are found. In [5], a participant A is considered honest when, in each possible context, she can always exculpate herself by a sequence of A -solo moves. Here we require that A is ready (i.e. some of her obligations are in the Weak Ready Do set) in all possible contexts, as in [3]. We conjecture that these two notions are equivalent. In [3] a type system has been proposed to safely over-approximate honesty. The type of a process P is a function which maps each variable to a *channel type*. These are behavioural types (in the form of Basic Parallel Processes) which essentially preserve the structure of P , by abstracting the actual prefixes as “non-blocking” and “possibly blocking”. The type system relies upon checking honesty for channel types, but no actual algorithm is given for such verification, hence type inference remains an open issue. In contrast, here we have directly implemented in Maude a verification algorithm for honesty, by model checking the abstract semantics in Sect. 5.

Acknowledgments. This work has been 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 “Behavioural Types for Reliable Large-Scale Software Systems” (BETTY).

References

1. Barbanera, F., de'Liguoro, U.: Two notions of sub-behaviour for session-based client/server systems. In: PDP (2010)
2. Bartoletti, M., Murgia, M., Scalas, A., Zunino, R.: Modelling and verifying contract-oriented systems in Maude. <http://tcs.unica.it/software/co2-maude>
3. Bartoletti, M., Scalas, A., Tuosto, E., Zunino, R.: Honesty by typing. In: Beyer, D., Boreale, M. (eds.) FORTE 2013 and FMOODS 2013. LNCS, vol. 7892, pp. 305–320. Springer, Heidelberg (2013)
4. Bartoletti, M., Tuosto, E., Zunino, R.: Contract-oriented computing in Co₂. *Sci. Ann. Comp. Sci.* **22**(1), 5–60 (2012)
5. Bartoletti, M., Tuosto, E., Zunino, R.: On the realizability of contracts in dishonest systems. In: Sirjani, M. (ed.) COORDINATION 2012. LNCS, vol. 7274, pp. 245–260. Springer, Heidelberg (2012)
6. Bartoletti, M., Zunino, R.: A calculus of contracting processes. In: LICS (2010)
7. Bruni, R., Corradini, A., Gadducci, F., Lluch Lafuente, A., Vandin, A.: Modelling and analyzing adaptive self-assembly strategies with maude. In: Durán, F. (ed.) WRLA 2012. LNCS, vol. 7571, pp. 118–138. Springer, Heidelberg (2012)
8. Castagna, G., Gesbert, N., Padovani, L.: A theory of contracts for web services. *ACM Trans. Program. Lang. Syst.* **31**(5), 1–61 (2009)
9. Clavel, M., Durán, F., Eker, S., Lincoln, P., Martí-Oliet, N., Meseguer, J., Quesada, J.F.: Maude: Specification and programming in rewriting logic. In: TCS (2001)
10. Honda, K., Vasconcelos, V.T., Kubo, M.: Language primitives and type discipline for structured communication-based programming. In: Hankin, C. (ed.) ESOP 1998. LNCS, vol. 1381, pp. 122–138. Springer, Heidelberg (1998)
11. Lange, J., Scalas, A.: Choreography synthesis as contract agreement. In: ICE (2013)
12. Meseguer, J.: Rewriting as a unified model of concurrency. In: Baeten, J.C.M., Klop, J.W. (eds.) CONCUR 1990. LNCS, vol. 458, pp. 384–400. Springer, Heidelberg (1990)
13. Meseguer, J.: Twenty years of rewriting logic. *JLAP* **81**(7–8), 721–781 (2012)
14. Neuhäuffer, M., Noll, T.: Abstraction and model checking of core Erlang programs in Maude. *ENTCS* **176**(4), 143–163 (2007)
15. Stehr, M.-O., Meseguer, J., Ölveczky, P.C.: Rewriting logic as a unifying framework for Petri Nets. In: Ehrig, H., Juhás, G., Padberg, J., Rozenberg, G. (eds.) APN 2001. LNCS, vol. 2128, pp. 250–303. Springer, Heidelberg (2001)
16. Thati, P., Sen, K., Martí-Oliet, N.: An executable specification of asynchronous pi-calculus semantics and may testing in Maude 2.0. In: ENTCS 71 (2002)
17. Verdejo, A., Martí-Oliet, N.: Implementing CCS in Maude 2. In: ENTCS 71 (2002)
18. Wirsing, M., Eckhardt, J., Mühlbauer, T., Meseguer, J.: Design and analysis of cloud-based architectures with KLAIM and maude. In: Durán, F. (ed.) WRLA 2012. LNCS, vol. 7571, pp. 54–82. Springer, Heidelberg (2012)
19. Şerbănuță, T.F., Roşu, G., Meseguer, J.: A rewriting logic approach to operational semantics. *Inf. Comput.* **207**(2), 305–340 (2009)