

Honesty by Typing

Massimo Bartoletti, Alceste Scalas, Emilio Tuosto, Roberto Zunino

► **To cite this version:**

Massimo Bartoletti, Alceste Scalas, Emilio Tuosto, Roberto Zunino. Honesty by Typing. Dirk Beyer; Michele Boreale. 15th International Conference on Formal Methods for Open Object-Based Distributed Systems (FMOODS) / 33th International Conference on Formal Techniques for Networked and Distributed Systems (FORTE), Jun 2013, Florence, Italy. Springer, Lecture Notes in Computer Science, LNCS-7892, pp.305-320, 2013, Formal Techniques for Distributed Systems. <10.1007/978-3-642-38592-6_21>. <hal-01515248>

HAL Id: hal-01515248

<https://hal.inria.fr/hal-01515248>

Submitted on 27 Apr 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Honesty by Typing

Massimo Bartoletti¹, Alceste Scalas¹, Emilio Tuosto², and Roberto Zunino³

¹ Università degli Studi di Cagliari, Italy — {bart, alceste.scalas}@unica.it

² University of Leicester, UK — emilio@mcs.le.ac.uk

³ Università degli Studi di Trento and COSBI, Italy — roberto.zunino@unitn.it

Abstract. We propose a type system for a calculus of contracting processes. Processes may stipulate contracts, and then either behave honestly, by keeping the promises made, or not. Type safety guarantees that a typeable process is *honest* — that is, the process abides by the contract it has stipulated in all possible contexts, even those containing dishonest adversaries.

1 Introduction

Most approaches to the formal specification of concurrent systems typically assume that components behave *honestly*, in that they always adhere to some agreed specification (e.g., a behavioural type), under the assumption that the static behaviour safely over-approximates the dynamic one. We argue that this assumption is unrealistic in scenarios where competition prevails against cooperation. Indeed, in a competitive scenario components may act selfishly, and diverge from the agreed specification.

We envision a *contract-oriented computing* paradigm [2], for the design of distributed components which use *contracts* to discipline their interaction. In CO₂ [2], a process may advertise contracts; a session is established processes advertising *compliant* contracts, similarly to other session-centric calculi. This session dictates the actions needed to realise their contracts to processes. A distinguished feature of CO₂ is that processes are not supposed to respect their contracts, nor are they bound to them by an enforcing mechanism.

More realistically, *dishonest* processes may avoid to perform some actions they have promised in their contracts. This may happen either intentionally, e.g. a malicious process trying to swindle the system, or unintentionally, e.g. because of some implementation bug (possibly exploited by some adversary). In both cases, the infrastructure can determine which process has caused the violation, and adequately punish it. A crucial problem is how to guarantee that a process will behave honestly, in *all possible contexts*. If such guarantee can be given, the process is protected both against unintentional bugs, and against (apparently honest) adversaries which try to make it sanctioned.

A negative result in [3] is that determining if a process is honest is an undecidable problem for a relevant class of contracts (introduced in [10], and refined in [11], for modelling WSDL and WSCL contracts). The problem is how to find a computable approximation of honesty, which implies the dynamic one.

Example. Let us consider an on-line store (participant A), which sells apples (a) and bottles of an expensive italian Brunello wine (b). Selling apples is quite easy: once an order is placed, A accepts it (with the feedback $\overline{\text{ok}}$) and waits for a payment ($\overline{\text{pay}}$) before shipping the goods ($\overline{\text{ship-a}}$). However, if expensive bottles of Brunello are ordered, the store is entitled to either decline the order (by answering $\overline{\text{no}}$), or accept it (and, as above, ship the item after the payment). Using external (+) and internal (\oplus) choice, the store contract can be modelled as

$$c = \text{a.}\overline{\text{ok}}.\overline{\text{pay}}.\overline{\text{ship-a}} + \text{b.}(\overline{\text{no}} \oplus \overline{\text{ok}}.\overline{\text{pay}}.\overline{\text{ship-b}})$$

External choice requires the other party to decide how to drive the contract evolution; in this case, the customer chooses between apples a and bottles b. Internal choice, instead, allows the advertising party to choose; in this case, the store selects either $\overline{\text{ok}}$ or $\overline{\text{no}}$.

Intuitively, contract compliance hinges on the duality between internal and external choices. Hence, to sell its goods, the store needs to find an agreement with a participant advertising a contract *compliant* with its contract c , such as:

$$d = \overline{\text{b}} . (\text{ok}.\overline{\text{pay}}.\overline{\text{ship-b}} + \text{no})$$

A buyer B advertising d wants to buy Brunello: she promises to select $\overline{\text{b}}$ (dual of b in the store contract), and then offers an external choice that lets the store choose between ok or no; in the first case, she promises to pay and wait for shipment.

In CO₂, the behaviour of each participant is a *process* capable of advertising contracts and executing the actions required to honour them. For instance, the store A can advertise its contract c by firing the prefix $\text{tell}_K \downarrow_x c$, where the index x in $\downarrow_x c$ is the name of a *channel* of A, and K is the name of an external broker, whom the contract is being advertised to. We shall not specify the behaviour of K, and just assume that when it finds a contract compliant with c , a session is established between A and another participant, say B. Technically, x is replaced with a fresh session name s . Participants A and B will then use such session to perform the actions required by their contracts.

A possible specification of the store A is e.g.:

$$\begin{aligned} P_M &= (x) (\text{tell}_K \downarrow_x c . (\text{do}_x \text{a} . X_M(x) + \text{do}_x \text{b} . X_M(x))) \\ X_M(x) &\stackrel{\text{def}}{=} \text{do}_x \overline{\text{ok}} . \text{do}_x \overline{\text{pay}} . \text{ask}_x \overline{\text{ship-a}}? . \text{do}_x \overline{\text{ship-a}} \end{aligned}$$

Here, A creates a private channel x , and advertises the contract c . Once a session starts, P_M can accept an order for a or b on x . This is modelled by the guards $\text{do}_x \text{a}$ and $\text{do}_x \text{b}$ of the choice operator + (not to be confused with + of contracts). In both cases, the process $X_M(x)$ is invoked. There, A accepts the transaction with ok, and waits for payment. Then A checks whether the contract requires to ship apples: if the query $\text{ask}_x \overline{\text{ship-a}}?$ passes, the goods are shipped. Otherwise, when the customer B orders Brunello, A maliciously gets stuck, and so B has paid for nothing. This store is *dishonest*: it does not respect its own contract c .

Consider now a non-malicious implementation of the store. Before accepting orders, the store requires an insurance to cover shipment damages. With the

contract $c_p = \overline{\text{payP}} . (\overline{\text{cover}} \oplus \overline{\text{cancel}})$, A promises to pay ($\overline{\text{payP}}$) and then choose between getting the coverage, or cancelling the request. The new store process is:

$$\begin{aligned}
P_N &= (x, y) (\text{tell}_C \downarrow_y c_p . \text{do}_y \overline{\text{payP}} . \text{tell}_K \downarrow_x c . \\
&\quad (\text{do}_x a . \text{do}_x \overline{\text{ok}} . X_N(x) + \text{do}_x b . Y_N(x, y))) \\
X_N(x) &\stackrel{\text{def}}{=} \text{do}_x \text{pay} . (\text{ask}_x \overline{\text{ship-a}}? . \text{do}_x \overline{\text{ship-a}} + \text{ask}_x \overline{\text{ship-b}}? . \text{do}_x \overline{\text{ship-b}}) \\
Y_N(x, y) &\stackrel{\text{def}}{=} \text{do}_y \overline{\text{cover}} . (\text{do}_x \overline{\text{ok}} . X_N(x) + \tau . \text{do}_x \overline{\text{no}})
\end{aligned}$$

The store A first requests an insurance by advertising c_p to an insurance company (say C); once C agrees, A pays the premium (on channel y), and then advertises c ; once an agreement with a customer B is reached, A waits for a or b orders. If apples are requested, A acknowledges ($\overline{\text{ok}}$) and invokes $X_N(x)$; there, A waits for payment, checks which good has to be shipped, and actually ships it. Otherwise, if Brunello is requested, $Y_N(x, y)$ is invoked: there, A requests the insurance coverage paid in advance; then, either the order is accepted and $X_N(x)$ is invoked for payment and shipment (as above), or the transaction is declined after an internal action τ (e.g. a wake up after a timeout).

This implementation is not malicious as the first attempt: it is keen to ship the goods, but it is not honest either due to the interaction between A and the insurance company. If C does not execute $\overline{\text{cover}}$, A gets stuck on $\text{do}_y \overline{\text{cover}}$, unable to honour c by providing the expected $\overline{\text{ok}}/\overline{\text{no}}$. Furthermore, A is dishonest w.r.t. c_p : the premium is paid in advance, but A may never perform $\text{do}_y \overline{\text{cover}}$ nor $\text{do}_y \overline{\text{cancel}}$ — e.g. if no agreement on c is found, or if the customer B is stuck, or if B simply chooses to buy apples. Thus, due to implementation naïveties, A may be blamed due to the unexpected (or malicious) behaviour of other participants.

Contributions. Our main contribution is a type discipline for statically ensuring when a CO_2 process is *honest*. The need for a static approximation is due to the fact that honesty is an undecidable property [3]. Our type system associates behavioural types to process channels. Checking honesty on these abstractions is decidable (Th. 1). We establish subject reduction (Th. 2) and progress (Th. 3), which are then used to prove type safety: typeable processes are honest (Th. 4).

Due to space constraints, we publish the proofs and further examples in a separate report [1].

2 A Calculus of Contracting Processes

The calculus CO_2 is parametric on the contract model, i.e. on the language of contracts. Contract models for CO_2 have been defined based e.g. on formulae of Propositional Contract Logic [4], and on CCS processes [2]. In this paper we focus on the two-party contracts of [11], as in [3]. Due to space limits, here we only give the main definitions, and refer the reader to [1] for the full details.

We assume a set of participants A, B, \dots , and a set of *atoms* a, b, \dots , that represent the actions performed by participants. We use an involution \bar{a} , as in

CCS, and assume a distinguished atom e , modelling a successfully terminated participant, such that $e = \bar{e}$.

A *unilateral contract* of [11] is a CCS-like process c which models the promised behaviour of a single participant. An internal sum $c = \bigoplus_{i \in I} \mathbf{a}_i ; c_i$ requires a participant (say, A) to choose one of the actions \mathbf{a}_i , do it, and then behave according to its continuation c_i . An external sum $c = \sum_{i \in I} \mathbf{a}_i . c_i$ requires A to offer a choice among all the branches; if \mathbf{a}_i is chosen by the other participant, then A must continue according to c_i . Finally, $c = \text{ready } \mathbf{a} . c'$ models an obligation of A to do \mathbf{a} , and then to proceed as c' . We denote (guarded) recursive contracts with $\text{rec } X . c$, and we let $E = \text{rec } X . e ; X$.

A *bilateral contract* $\gamma = A \text{ says } c \mid B \text{ says } d$ combines the contracts of two participants. Its behaviour is given in [3] as a LTS. Its main rules regulate the interaction between the internal choices of A , and the external choices of B :

$$A \text{ says } (\bar{\mathbf{a}} ; c \oplus c') \mid B \text{ says } (\mathbf{a} . d + d') \xrightarrow{A \text{ says } \mathbf{a}} A \text{ says } c \mid B \text{ says } \text{ready } \mathbf{a} . d \quad (1)$$

$$A \text{ says } c \mid B \text{ says } \text{ready } \mathbf{a} . d \xrightarrow{B \text{ says } \mathbf{a}} A \text{ says } c \mid B \text{ says } d \quad (2)$$

By (1), if A chooses branch $\bar{\mathbf{a}}$ in her internal sum, then B is committed to the branch \mathbf{a} in his external sum. We mark the selected branch with *ready* \mathbf{a} , and discard the others. This enables (2) where B takes the marked branch. When an internal choice is not matched by an action in the external choice of the partner, the contract gets stuck. Intuitively, two contracts c, d are *compliant*, $c \bowtie d$ in symbols, when this situation cannot occur. The formal definition of compliance builds upon that of *ready sets*. For a contract c , the ready sets $RS(c)$ are:

$$\begin{array}{ll} \{\{\text{ready } \mathbf{a}\}\}, & \text{if } c = \text{ready } \mathbf{a} . c' \\ \{\{\mathbf{a}_i\} \mid i \in I\}, & \text{if } c = \bigoplus_{i \in I} \mathbf{a}_i ; c_i \text{ and } I \neq \emptyset \end{array} \quad \begin{array}{ll} RS(c'), & \text{if } c = \text{rec } X . c' \\ \{\{\mathbf{a}_i \mid i \in I\}\}, & \text{if } c = \sum_{i \in I} \mathbf{a}_i . c_i \end{array}$$

Then, the relation \bowtie on contracts is the largest relation preserved by contract transitions such that, whenever $c \bowtie d$,

$$\forall \mathcal{X} \in RS(c), \mathcal{Y} \in RS(d). (\{\bar{\mathbf{a}} \mid \mathbf{a} \in \mathcal{X}\} \cap \mathcal{Y} \neq \emptyset \text{ or } \exists \mathbf{a} . \text{ready } \mathbf{a} \in (\mathcal{X} \cup \mathcal{Y}) \setminus (\mathcal{X} \cap \mathcal{Y}))$$

We now briefly review CO_2 . Let \mathcal{V} and \mathcal{N} be disjoint sets of, respectively, *session variables* x, y, \dots and *session names* s, t, \dots . Let u, v, \dots range over $\mathcal{V} \cup \mathcal{N}$. *Systems* S , *processes* P , *prefixes* π , and *latent contracts* K are defined below.

Definition 1 (CO₂ syntax). *The syntax of CO₂ is given by:*

$$\begin{array}{ll} P ::= \sum_i \pi_i . P_i \mid P \mid P \mid (\bar{u})P \mid X(\bar{u}) & \pi ::= \tau \mid \text{tell}_{\mathbf{A}} \downarrow_u c \mid \text{fuse} \mid \text{do}_u \mathbf{a} \mid \text{ask}_u \phi \\ K ::= \downarrow_u A \text{ says } c \mid K \mid K & S ::= \mathbf{0} \mid A[P] \mid A[K] \mid s[\gamma] \mid S \mid S \mid (\bar{u})S \end{array}$$

Processes specify the behaviour of participants. A process can be a prefix-guarded finite sum $\sum_i \pi_i . P_i$, a parallel composition $P \mid Q$, a delimited process $(\bar{u})P$, or a constant $X(\bar{u})$. We write $\mathbf{0}$ for $\sum_{\emptyset} P$, and $\pi_1 . Q_1 + P$ for $\sum_{i \in I \cup \{1\}} \pi_i . Q_i$ provided that $P = \sum_{i \in I} \pi_i . Q_i$ and $1 \notin I$. We omit trailing occurrences of $\mathbf{0}$.

We stipulate that each X has a unique defining equation $X(u_1, \dots, u_j) \stackrel{\text{def}}{=} P$ s.t. $\text{fv}(P) \subseteq \{u_1, \dots, u_j\} \subseteq \mathcal{V}$, and each constant occurring in P is prefix-guarded.

Prefixes include the silent action τ , contract advertisement $\text{tell}_A \downarrow_u c$, contract stipulation fuse , action execution $\text{do}_u a$, and contract query $\text{ask}_u \phi$. In each prefix $\pi \neq \tau$, the identifier u refers to the target session involved in the execution of π . As in [3], we leave the syntax of *observables* ϕ unspecified.

A *latent contract* $\downarrow_x A$ says c represents a contract c advertised by A but not stipulated yet. The variable x will be instantiated to a fresh session name upon stipulation. K simply stands for the parallel composition of latent contracts.

A system is composed of participants $A[P]$, sessions $s[\gamma]$, sets of latent contracts advertised to A (denoted by $A[K]$), and delimited systems $(\bar{u})S$. Delimitation (\bar{u}) binds session variables and names, both in processes and systems. Free variables and names are defined as usual, and denoted by $\text{fv}(\cdot)$ and $\text{fn}(\cdot)$. A system/process is *closed* when it has no free variables. Each participant may have at most one process, i.e. we forbid systems of the form $A[P] \mid A[Q]$. We say that S is *A-free* when it does not contain the participant $A[P]$, nor latent contracts of A , nor sessions with A 's contracts. Note that sessions cannot contain latent contracts.

The semantics of CO_2 is formalised by a reduction relation on systems (Def. 2). This relies on a standard structural congruence relation — of which we just point out that $(\bar{u})A[(\bar{v})P] \equiv (\bar{u}\bar{v})A[P]$ allows to move delimitations between systems and processes, while $A[K] \mid A[K'] \equiv A[K \mid K']$ allows A to collect latent contracts.

In order to define honesty in § 3, here we decorate transitions with labels, by writing $\xrightarrow{A : \pi, \sigma}$ for a reduction where participant A fires prefix π . Also, σ is a substitution which accounts for the instantiation of session variables upon a *fuse*.

Definition 2 (CO₂ semantics). *The relation $\xrightarrow{A : \pi, \sigma}$ between systems (considered up-to structural congruence \equiv) is the smallest relation closed under the rules of Fig. 1. The relation $K \triangleright^\sigma \gamma$ holds iff (i) K has the form $\downarrow_x A$ says $c \mid \downarrow_y B$ says d , (ii) $c \bowtie d$, (iii) $\gamma = A$ says $c \mid B$ says d , and (iv) $\sigma = \{s/x, y\}$ maps $x, y \in \mathcal{V}$ to $s \in \mathcal{N}$. The substitution $\sigma_{\neq u}$ in rule [DEL2] is defined as $\sigma(v)$ for all $v \neq u$, and it is undefined on u .*

The rules in Fig. 1 are a minor variation of those presented in [3]. Their intuitive meaning is sketched in the introductory example (Sec. 1): [TELL] advertises a contract c , [FUSE] creates a new session s upon contractual compliance, [DO] performs a contractual action, [ASK] blocks until session s satisfies observable ϕ .

Example 1. A possible execution of $S = A[(x)X(x)] \mid B[(y)Y(y)] \mid C[\text{fuse}]$, where $X(x) \stackrel{\text{def}}{=} \text{tell}_C \downarrow_x (a; E) . \text{do}_x a$ and $Y(y) \stackrel{\text{def}}{=} \text{tell}_C \downarrow_y (\bar{a} . E) . \text{do}_y \bar{a}$, is:

$$\begin{aligned}
S & \xrightarrow{B : \text{tell}_C \downarrow_y \bar{a}, \emptyset} A[(x)X(x)] \mid C[\text{fuse}] \mid (y) (B[\text{do}_y \bar{a}] \mid C[\downarrow_y B \text{ says } \bar{a} . E]) \\
& \xrightarrow{A : \text{tell}_C \downarrow_x a, \emptyset} (x) (A[\text{do}_x a] \mid (y) (B[\text{do}_y \bar{a}] \mid C[\text{fuse}] \mid C[\downarrow_x A \text{ says } a; E \mid \downarrow_y B \text{ says } \bar{a} . E])) \\
& \xrightarrow{C : \text{fuse}, \emptyset} (s) (A[\text{do}_s a] \mid (y) (B[\text{do}_s \bar{a}] \mid C[0] \mid s[A \text{ says } a; E \mid B \text{ says } \bar{a} . E])) \\
& \xrightarrow{A : \text{do}_s a, \emptyset} (s) (A[0] \mid B[\text{do}_s \bar{a}] \mid s[A \text{ says } E \mid B \text{ says ready } \bar{a} . E])
\end{aligned}$$

$$\begin{array}{c}
\frac{A[\tau.P + P' \mid Q] \xrightarrow{A: \tau, \emptyset} A[P \mid Q] \quad [TAU] \quad \frac{S \xrightarrow{A: \pi, \sigma} S' \quad \text{ran } \sigma \cap \text{fn}(S'') = \emptyset}{S \mid S'' \xrightarrow{A: \pi, \sigma} S' \mid S'' \sigma} \quad [PAR]}{A[\text{tell}_B \downarrow_u c.P + P' \mid Q] \xrightarrow{A: \text{tell}_B \downarrow_u c, \emptyset} A[P \mid Q] \mid B[\downarrow_u A \text{ says } c]} \quad [TELL]} \\
\frac{\frac{K \triangleright^\sigma \gamma \quad \text{ran } \sigma = \{s\} \quad s \text{ fresh}}{A[\text{fuse}.P + P' \mid Q] \mid A[K] \xrightarrow{A: \text{fuse}, \sigma} A[P \mid Q]\sigma \mid s[\gamma]} \quad [FUSE]}{\frac{\gamma \xrightarrow{A \text{ says } a} \gamma'}{A[\text{do}_s a.P + P' \mid Q] \mid s[\gamma] \xrightarrow{A: \text{do}_s a, \emptyset} A[P \mid Q] \mid s[\gamma']} \quad [DO]}{\frac{\gamma \vdash \phi}{A[\text{ask}_s \phi.P + P' \mid Q] \mid s[\gamma] \xrightarrow{A: \text{ask}_s \phi, \emptyset} A[P \mid Q] \mid s[\gamma]} \quad [ASK] \quad \frac{S \xrightarrow{A: \pi, \{s/x\}} S'}{(x)S \xrightarrow{A: \pi, \emptyset} (s)S'} \quad [DELT1]}{S \xrightarrow{A: \pi, \sigma} S' \quad \frac{u \notin \text{ran } \sigma \quad \sigma \neq u \neq \emptyset}{(u)S \xrightarrow{A: \pi, \sigma \neq u} (u)S'} \quad [DELT2] \quad \frac{X(\vec{u}) \stackrel{\text{def}}{=} P \quad A[P\{\vec{v}/\vec{u}\} \mid Q] \mid S \xrightarrow{A: \pi, \sigma} S'}{A[X(\vec{v}) \mid Q] \mid S \xrightarrow{A: \pi, \sigma} S'} \quad [DEF]}
\end{array}$$

Fig. 1. Reduction semantics of CO₂.

3 On Honesty

A participant A is honest when she *realizes* every contract she advertises, in every session she may be engaged in. If a system S contains a session s with a contract c advertised by A , such as $A[P \mid s[A \text{ says } c \mid \dots] \mid \dots]$, then A must realize c , even in a system populated by adversaries who play to cheat her. To realize c , A must be “ready” to behave according to c . For instance, if $A[P]$ has advertised a contract c with an internal choice $c_i = a \oplus b$, then P must be ready to do *at least one* of the actions a, b . Instead, if c is an external choice $c_e = a + b$, then P must be ready to do *both* the actions a and b .

Realizability requires the *readiness* property to be preserved by all transitions of S . In other words, in any reduct of S containing a reduct P' of P and a reduct c' of c , the process P' must be ready for c' . To formalise when “ P is ready for c ” we inspect the ready sets $RS(c)$, which reveal whether c is exposing an internal or an external choice. At the process level, we consider the reachable actions in P .

Example 2. Let $c_i = a \oplus b$, and let $c_e = a + b$. Then, $RS(c_i) = \{\{a\}, \{b\}\}$, and $RS(c_e) = \{\{a, b\}\}$. Consider now the following processes:

- $P_0 = \text{do}_s a$ is ready for c_i , because $\{a\} \in RS(c_i)$ and $\text{do}_s a$ is enabled in P_0 . Instead, P_0 is *not* ready for c_e , since the ready set $\{a, b\}$ of c_e also contains b , which is not enabled in P_0 .
- $P_1 = \text{do}_s a + \text{do}_s b + \text{do}_s z$ is ready for both c_i and c_e . Indeed, P_1 enables $\text{do}_s a$ and $\text{do}_s b$, thus covering the ready sets of c_i and c_e . The branch $\text{do}_s z$ is immaterial: rule [DO] blocks actions not expected by the contract.
- $P_2 = \tau.\text{do}_s a + \tau.\text{do}_s b$ is ready for c_i , because whatever branch is taken by P_2 , it leads to an unguarded action which covers one of the ready sets

in c_i . Instead, P_2 is *not* ready for c_e , because after one of the two branches is chosen, one of the two actions expected by c_e is no longer available.

- $P_3 = \text{do}_t w.\text{do}_s a + \text{do}_t z.\text{do}_s b$ is a bit more complex than the above cases. Readiness w.r.t. c_i depends on the context. If the context eventually enables one of the do_t , then either $\text{do}_s a$ or $\text{do}_s b$ will be enabled, hence P_3 is ready for c_i . Otherwise, P_3 is stuck, hence it is not ready for c_i . Notice that P_3 is not ready for c_e , regardless of the context.

To formalise readiness, we start by defining the set $RD_u^A(S)$ (for “Ready Do”), which collects all the atoms with an unguarded action do_u of A in a system S .

Definition 3 (Ready do). For all S, A and u , we define the set $RD_u^A(S)$ as:

$$RD_u^A(S) = \{a \mid \exists \vec{v}, P, P', Q, S' . S \equiv (\vec{v}) (A[\text{do}_u a.P + P' \mid Q] \mid S') \wedge u \notin \vec{v}\}$$

As seen for P_2 and P_3 , readiness may also hold when the actions expected in the contract ready sets are not immediately enabled in the process. To check if $A[P]$ is ready for s (in a system S), we need to consider all actions which (1) are exposed in P after some steps, taken by P itself or by the context, and (2) are not preceded by other do_s performed by A . These actions form the set $WRD_s^A(S)$.

Definition 4 (Weak ready do). We define the set of atoms $WRD_u^A(S)$ as:

$$WRD_u^A(S) = \{a \mid \exists S' : S \xrightarrow{\neq(A: \text{do}_u)}^* S' \text{ and } a \in RD_u^A(S')\}$$

where $S \xrightarrow{\neq(A: \text{do}_u)} S'$ iff $\exists B, \pi, \sigma. S \xrightarrow{B: \pi, \sigma} S' \wedge (B \neq A \vee \forall a. \pi \neq \text{do}_u a)$.

Example 3. For $S = A[\text{do}_x \bar{a} . \text{do}_y b + \tau . \text{do}_y a . \text{do}_y c \mid (x) \text{do}_x \bar{b}]$, we have:

$$WRD_x^A(S) = \{\bar{a}\} = RD_x^A(S) \quad WRD_y^A(S) = \{a\} \supseteq RD_y^A(S) = \emptyset$$

On channel y , the action a is weakly reachable through its τ prefix. Action b is *not* weakly reachable, because guarded by a stuck do_x . Action c is *not* weakly reachable as well, because preceded by another do on the same channel.

Example 4. Let P_3 as in Ex. 2, and consider the following system:

$$S = A[P_3] \mid B[\tau . \text{do}_s \bar{a} . \text{do}_s \bar{b}] \mid C[\text{do}_t \bar{w} + \text{do}_t \bar{z} + \tau] \\ \mid s[A \text{ says } a + b \mid B \text{ says } \bar{a} \oplus \bar{b}] \mid t[A \text{ says } w + z \mid C \text{ says } \bar{w} \oplus \bar{z}]$$

In session t , A is immediately ready to do either w or z , so her ready do set coincides with her weak ready do set in t . The same for C , with the dual atoms \bar{w} and \bar{z} . Thus:

$$WRD_t^A(S) = RD_t^A(S) = \{w, z\} \quad WRD_t^C(S) = RD_t^C(S) = \{\bar{w}, \bar{z}\}$$

In session s , the ready do sets of both A and B are empty, because their actions are not immediately enabled. Before they can be reached, the whole system S

must first reduce, either with the contribution of **C** on session t (in the case of **A**), or through a τ action (in the case of **B**). These reductions fall within the definition of their *weak* ready do sets, which are accordingly non-empty.

$$WRD_s^B(S) = \{\bar{a}\} \supseteq RD_s^B(S) = \emptyset \quad WRD_s^A(S) = \{a, b\} \supseteq RD_s^A(S) = \emptyset$$

Notice that $\bar{b} \notin WRD_s^B(S)$: in fact, \bar{b} is guarded by $\text{do}_s \bar{a}$. Also, if **C** chooses to perform τ , then the actions in $WRD_s^A(S)$ would not be reached. Indeed, Def. 4 requires that each element in the set becomes reachable at the end of a suitable reduction trace — but it does not prevent S from reducing along other paths.

A participant **A** is ready in a system S with a session $s[\mathbf{A} \text{ says } c \mid \dots]$ iff **A** is (weakly) ready to do all the actions in some ready set of c . Note that **A** is vacuously ready in systems not containing sessions with **A**'s contracts.

Definition 5 (Honesty). *We say that **A** is ready in S iff, whenever $S \equiv (\bar{u})S'$ for some \bar{u} and $S' = s[\mathbf{A} \text{ says } c \mid \dots] \mid S_0$,*

$$\exists \mathcal{X} \in RS(c) . \forall a \neq e . (a \in \mathcal{X} \vee \text{ready } a \in \mathcal{X} \implies a \in WRD_s^A(S'))$$

$A[P]$ is honest iff \forall **A**-free S and $\forall S'$ such that $A[P] \mid S \rightarrow^* S'$, **A** is ready in S' .

The **A**-freeness requirement in Def. 5 is used just to rule out those systems already carrying stipulated or latent contracts of **A** outside $A[P]$, e.g. $A[P] \mid B[\downarrow_x \mathbf{A} \text{ says } \bar{p}a\bar{y} \mid \dots]$. In the absence of **A**-freeness, the system could trivially make $A[P]$ dishonest.

Example 5. In the system below, **A** might look honest, although she is not.

$$\begin{aligned} S &\stackrel{\text{def}}{=} A[(x, y) (P_A \mid \text{fuse} \mid \text{fuse})] \mid B[P_B] \mid C[P_C] \\ P_A &\stackrel{\text{def}}{=} \text{tell}_A(\downarrow_x a . E) . \text{tell}_A(\downarrow_y b ; E) . \text{do}_x a . \text{do}_y b \\ P_B &\stackrel{\text{def}}{=} (z) (\text{tell}_A(\downarrow_z \bar{b} . E) . \text{do}_z \bar{b}) \quad P_C \stackrel{\text{def}}{=} (w) (\text{tell}_A(\downarrow_w \bar{a} ; E) . \mathbf{0}) \end{aligned}$$

In fact, if we reduce S by performing all the **tell** and **fuse** actions, we obtain:

$$\begin{aligned} S' = (s, t) (& A[\text{do}_t a . \text{do}_s b] \mid B[\text{do}_s \bar{b}] \mid C[\mathbf{0}] \mid \\ & t[A \text{ says } a . E \mid C \text{ says } \bar{a} ; E] \mid s[A \text{ says } b ; E \mid B \text{ says } \bar{b} . E]) \end{aligned}$$

Here, S' cannot reduce further: **A** is stuck, waiting for \bar{a} from **C**, which (dishonestly) avoids to do the required internal choice. So, **A** is dishonest, because she does not perform the promised **b**. Formally, **A** is dishonest because $RS(b; E) = \{\{b\}\}$, but $b \notin WRD_s^A(S')$. Thus, **A** is not ready in S' , hence not honest in S .

Our definition of honesty subsumes a *fair* scheduler, eventually allowing participants to fire persistently (weakly) enabled **do** actions. For instance, let $c = a \oplus b$, and let:

$$P \stackrel{\text{def}}{=} (x) (\text{tell}_A \downarrow_x c . \text{fuse} . X(x)) \quad \text{where } X(x) \stackrel{\text{def}}{=} \tau . X(x) + \tau . \text{do}_x a + \tau . \text{do}_x b$$

Let $S = A[P] \mid S_0$, and assume that the **fuse** in P passes. Under an unfair scheduler, **A** could always take the first branch in X , while neglecting the others. Intuitively, this would make **A** not respect her contract, which expects **a** or **b**. However, a fair scheduler will eventually choose one of the other branches.

4 A Type System for CO₂

We now introduce a type system for CO₂. The main result is *type safety* (established in Th. 4), which guarantees that typeable participants are honest.

The type of a process P is a function f , which maps each channel to a *channel type*. Channel types are behavioural types which essentially preserve the structure of P , while abstracting the actual prefixes and delimitations. Mainly, the prefixes of channel types distinguish between nonblocking and possibly blocking actions.

4.1 Channel types

Channel types are Basic Parallel Processes (BPPs [14]) with standard semantics. Their prefixes can be atoms ($\mathbf{a}, \mathbf{b}, \dots$), contract advertisement actions ($\langle c \rangle$), nonblocking (τ), possibly blocking ($\tau?$), and conditional (τ_ϕ) silent actions.

Definition 6 (Channel types). Channel types T and prefixes α are:

$$T ::= \mathbf{0} \mid \alpha . T \mid T + T \mid T \mid T \mid \text{rec } X . T \mid X \quad \alpha ::= \mathbf{a} \mid \tau \mid \tau? \mid \tau_\phi \mid \langle c \rangle$$

Example 6. Let $P = \text{tell}_{\mathbb{B}} \downarrow_x c_i \mid (\text{tell}_{\mathbb{B}} \downarrow_y d . \text{do}_x \bar{\mathbf{a}})$, where $c_i = \bar{\mathbf{a}} \oplus \bar{\mathbf{b}}$, and d is immaterial. We anticipate that the channel types associated by our type system to P on channels x and y are, respectively, $T_x = \langle c_i \rangle \mid \tau . \bar{\mathbf{a}}$, and $T_y = \tau \mid \langle d \rangle . \tau?$. Note that the advertisement of $\downarrow_x c_i$ is recorded in T_x , while that of $\downarrow_y d$ is abstracted there as a τ . Instead, the $\tau?$ in T_y represents the fact that $\text{do}_x \bar{\mathbf{a}}$ is not visible from channel y , and may potentially block.

The execution of CO₂ systems relies both on processes and contracts. Thus, we use an abstraction of the contract/process interplay to define abstract processes.

Definition 7 (Abstract processes). An abstract process is a pair (C, T) or (c, T) , where C is a set of contracts, c is a contract, and T is a channel type.

An abstract process (C, T) represents a process abstracted by T on some channel x , after the contracts in C have been *advertised*. Instead, (c, T) represents a process abstracted by T on channel x , after the contract c has been *stipulated*.

The semantics of abstract processes is given in Fig. 2. The set C grows when a channel type T in (C, T) performs a transition with label $\langle c \rangle$. After a contract $c \in C$ has been stipulated, the set is reduced to c . A label \mathbf{a} models a $\text{do } \mathbf{a}$ action performed by T , while rule ctx models an (unknown) action performed by the context. Further advertisements after contract stipulation are neglected.

The relation $\rightarrow_{\#}$ abstracts the contract semantics \rightarrow , by considering only the contract advertised by P (instead of the whole bilateral contract). We leave the relation $\rightarrow_{\#}$ unspecified (see [3] for a possible instantiation), and we just require that $\rightarrow_{\#}$ is decidable, and for all $\gamma = \mathbf{A} \text{ says } c \mid \mathbf{B} \text{ says } d$ such that $c \bowtie d$,

$$\begin{aligned} \gamma \xrightarrow{\mathbf{A} \text{ says } \mathbf{a}} \mathbf{A} \text{ says } c' \mid \mathbf{B} \text{ says } d' &\Longrightarrow c \xrightarrow{\mathbf{a}}_{\#} c' \\ \gamma \xrightarrow{\mathbf{B} \text{ says } \mathbf{b}} \mathbf{A} \text{ says } c' \mid \mathbf{B} \text{ says } d' &\Longrightarrow c \xrightarrow{ctx}_{\#} c' \end{aligned}$$

$$\begin{array}{c}
\frac{T \xrightarrow{\langle c \rangle} T'}{(C, T) \rightarrow (C \cup \{c\}, T')} \quad \frac{T \xrightarrow{\langle d \rangle} T'}{(c, T) \rightarrow (c, T')} \quad \frac{c \in C}{(C, T) \rightarrow (c, T)} \quad \frac{c \xrightarrow{ctx} \# c'}{(c, T) \rightarrow (c', T)} \\
\frac{T \xrightarrow{\alpha} T' \quad \alpha \in \{\tau, \tau?, \tau_\phi\}}{(C, T) \rightarrow (C, T')} \quad \frac{c \xrightarrow{a} \# c' \quad T \xrightarrow{a} T'}{(c, T) \rightarrow (c', T')} \quad \frac{T \xrightarrow{\alpha} T' \quad \alpha \in \{\tau, \tau?, \tau_\phi\}}{(c, T) \rightarrow (c, T')} \\
\frac{T \xrightarrow{a} T'}{T \xrightarrow{a} T'} \quad \frac{T \xrightarrow{\tau} T'' \xrightarrow{a} T'}{T \xrightarrow{a} T'} \quad \frac{T \xrightarrow{\langle d \rangle} T'' \xrightarrow{a} T'}{T \xrightarrow{a} T'} \quad \frac{T \xrightarrow{\tau_\phi} T'' \xrightarrow{a} T' \quad c \vdash \#^A \phi}{T \xrightarrow{a} T'}
\end{array}$$

Fig. 2. Abstract processes semantics and channel type semantics.

We now introduce the abstract counterpart of the dynamic notion of honesty in § 3. We shall follow the path outlined for concrete processes: first we define when a channel type T is “ready for a contract”, and then when T is honest.

Weak transitions abstract the weak ready do. They are defined in Fig. 2 as a labelled relation \xrightarrow{a}^A_c (simply written as \xrightarrow{a} when unambiguous). The first two rules are standard: they just collapse the τ actions as usual. The third rule also collapses contract advertisement actions, which are nonblocking as well. Possibly blocking actions $\tau?$ are *not* collapsed, while τ_ϕ (which abstract $\text{ask}_u \phi$ prefixes) are dealt with the last rule: they abstract the CO_2 prefix $\text{ask}_u \phi$, and they are collapsed only if such ask is nonblocking. The relation $\vdash \#^A$ safely (under-) approximates this condition. We leave $\vdash \#^A$ unspecified (just like \vdash in § 2), and we only require that it respects the constraint in Def. 8 below. Unlike in the concrete case, the context is immaterial in determining weak transitions.

Definition 8 (Abstract observability). *We write $c \vdash \#^A \phi$ for any decidable relation satisfying: $c \vdash \#^A \phi \implies \forall B. \forall d. (c \bowtie d \implies A \text{ says } c \mid B \text{ says } d \vdash \phi)$.*

Below we relate the weak transitions of a channel type with the ready sets of a contract, similarly to what we did in Def. 5 for CO_2 processes. Weak transitions under-approximate the weak ready do set. Thus, if an abstract process is honest then also the concrete one will be such (while the *vice versa* is not always true).

Honesty of abstract processes is defined similarly to Def. 5. In order to be honest, a process must keep itself (abstractly) ready upon transitions. Readiness must be checked against all the contracts that may be stipulated along the abstract process reductions.

Definition 9 (Abstract honesty). *We say that channel type T is abstractly ready for contract c iff $\exists \mathcal{X} \in RS(c). \forall a \neq e. ((a \in \mathcal{X} \vee \text{ready } a \in \mathcal{X}) \implies T \xrightarrow{a})$. An abstract process $(-, T)$ is honest iff, whenever $(-, T) \rightarrow^* (c, T')$, then T' is ready for c . A channel type T is honest iff (\emptyset, T) is honest.*

Example 7. Let $T_x = \langle c_i \rangle \mid \tau. a$, and recall the contract $c_i = a \oplus b$ from Ex. 2. To prove T_x is honest, we examine all the reducts of the abstract process (\emptyset, T_x) to check for readiness. We have the following cases:

1. (\emptyset, T_x) . Nothing to check, because no contracts have been advertised yet.
2. $(\emptyset, \langle c_i \mid \mathbf{a} \rangle)$. Similar to the previous case.
3. $(\{c_i\}, \tau . \mathbf{a})$. Nothing to check, since no contracts have been stipulated yet.
4. $(\{c_i\}, \mathbf{a})$. Similar to the previous case.
5. $(c_i, \tau . \mathbf{a})$. Here $\tau . \mathbf{a}$ is ready for c_i , as $\{\mathbf{a}\} \in RS(c_i) = \{\{\mathbf{a}\}, \{\mathbf{b}\}\}$ and $\tau . \mathbf{a} \xrightarrow{\mathbf{a}}$.
6. (c_i, \mathbf{a}) . We have that \mathbf{a} is ready for c_i , similarly to the previous case.
7. $(E, \mathbf{0})$. We have that $\mathbf{0}$ is vacuously ready for E .

Th. 1 below establishes that checking the honesty of a channel type T is decidable. Indeed, abstract readiness and abstract dishonesty are reachability properties. Abstract processes are the product of a finite state system (C and c only admit finitely many states), and a BPP T . This product can be modelled as a Petri net. Decidability follows since reachability is decidable for Petri nets [17].

Theorem 1. *Abstract honesty is decidable.*

4.2 Process types

A CO_2 process type associates session names/variables to channel types, thus abstracting the behaviour of a process on all channels. Here, we introduce a “dummy” channel $*$ $\notin \mathcal{N} \cup \mathcal{V}$, for collecting type information about unused channels. Formally, a process type is a function f from $\mathcal{N} \cup \mathcal{V} \cup \{*\}$ to channel types.

Our type system abstracts concrete prefixes of CO_2 processes as actions of channel types. We observe the behaviour of a process P on each channel, say u . When P performs an action on one of its channels, say v , we have two cases:

$v \neq u$: we will only observe a silent action, either nonblocking (τ) or blocking ($\tau?$), depending on the concrete prefix fired.

$v = u$: we may observe more information, depending on the concrete prefix.

For instance, if P advertises a contract c with a $\text{tell}_{\downarrow v} c$, then the action $\langle c \rangle$ will be visible if $v = u$, while we shall just observe a τ if $v \neq u$ (because tell is nonblocking). Similarly, if P performs $\text{do}_v \mathbf{a}$ we shall observe the action \mathbf{a} if $v = u$ and $\tau?$ if $v \neq u$ (because do is blocking). Finally, if P executes a query $\text{ask}_u \phi$ we shall observe the conditional silent action τ_ϕ if $v = u$ and $\tau?$ otherwise.

Definition 10 (Prefix abstraction). *For all $u \in \mathcal{N} \cup \mathcal{V} \cup \{*\}$, we define the mapping $[\cdot]_u$ from CO_2 prefixes to channel type prefixes as follows:*

$$\begin{aligned}
[\tau]_u &= \tau & [\text{fuse}]_u &= \tau? & [\text{tell}_{\downarrow v} c]_u &= \text{if } v = u \text{ then } \langle c \rangle \text{ else } \tau \\
[\text{do}_v \mathbf{a}]_u &= \text{if } v = u \text{ then } \mathbf{a} \text{ else } \tau? & [\text{ask}_v \phi]_u &= \text{if } v = u \text{ then } \tau_\phi \text{ else } \tau?
\end{aligned}$$

The typing judgments for processes have the form $\Gamma \vdash P : f$, where Γ is a typing environment, giving types to processes of the form $X(\vec{v})$.

Definition 11 (Typing environment). *A typing environment Γ is a partial function associating process types to constants $X(\vec{v})$.*

$$\begin{array}{c}
\frac{\Gamma \vdash P_i : f_i \quad \forall i \in I}{\Gamma \vdash \sum_{i \in I} \pi_i \cdot P_i : \lambda u . \sum_{i \in I} [\pi_i]_u \cdot f_i(u)} \text{[T-SUM]} \quad \frac{\Gamma \vdash P : f \quad \Gamma \vdash Q : g}{\Gamma \vdash P \mid Q : \lambda u . f(u) \mid g(u)} \text{[T-PAR]} \\
\frac{X(\vec{u}) \stackrel{\text{def}}{=} P \quad \Gamma \{f/X(\vec{v})\} \vdash P \{\vec{v}/\vec{u}\} : f}{\Gamma \vdash X(\vec{v}) : f} \text{[T-DEF]} \quad \frac{\Gamma(X(\vec{v})) = f}{\Gamma \vdash X(\vec{v}) : f} \text{[T-VAR]} \\
\frac{\Gamma_{\neq u} \vdash P : f \quad f(u) \text{ honest}}{\Gamma \vdash (u)P : f \{f^*/u\}} \text{[T-DEL]} \quad \text{where } \Gamma_{\neq \vec{v}}(Y(\vec{w})) = \begin{cases} \Gamma(Y(\vec{w})) & \text{if } \vec{w} \cap \vec{v} = \emptyset \\ \text{undefined} & \text{otherwise} \end{cases}
\end{array}$$

Fig. 3. Typing rules for processes.

In Fig. 3 we introduce the typing rules for CO₂ processes. Rule [T-SUM] abstracts the prefixes which guard the branches of a summation, according to Def. 10. The resulting process type is expressed through the usual λ -notation. The type of a parallel composition is the pointwise parallel composition of the component types (rule [T-PAR]). Rules [T-DEF] and [T-VAR] are mostly standard. Rule [T-VAR] retrieves the type of a process variable from the typing environment, which is populated by [T-DEF]. The rule for typing delimitations ([T-DEL]) is worth some comments. Assume that P is typed with f . Since u is not free in $(u)P$, the actions on channel u must not be observable in the typing of $(u)P$. To do that, in the typing of $(u)P$ we discard the information on u , by replacing it with the typing information on the “dummy” channel $*$. However, since this might hide a dishonest behaviour on u , the rule also checks that $f(u)$ is honest. Moreover, if the environment Γ has typing information on channel u , this cannot be used while typing P . The typing environment $\Gamma_{\neq u}$, which discards the information on u , is used to this purpose.

Example 8. Recall process $P_2 = \tau \cdot \text{do}_s \mathbf{a} + \tau \cdot \text{do}_s \mathbf{b}$ from Ex. 2. Its typing derivation is:

$$\frac{\frac{\vdash \text{do}_s \mathbf{a} : \lambda u . [\text{do}_s \mathbf{a}]_u = f_1}{\vdash \text{do}_s \mathbf{a} : \lambda u . [\text{do}_s \mathbf{a}]_u = f_1} \text{[T-SUM]} \quad \frac{\vdash \text{do}_s \mathbf{b} : \lambda u . [\text{do}_s \mathbf{b}]_u = f_2}{\vdash \text{do}_s \mathbf{b} : \lambda u . [\text{do}_s \mathbf{b}]_u = f_2} \text{[T-SUM]}}{\vdash P_2 : f = \lambda u . [\tau]_u \cdot f_1(u) + [\tau]_u \cdot f_2(u)} \text{[T-SUM]}$$

We have $f(s) = \tau \cdot \mathbf{a} + \tau \cdot \mathbf{b}$, and for all $u \neq s$, $f(u) = f(*) = \tau \cdot \tau_? + \tau \cdot \tau_?$. In other words, the process type f performs some visible actions when observed at channel s , while remaining “silent” on other channels.

The type system assigns the same type (up-to structural congruence) to all non-free session names/variables, and to $*$; such type may only contain τ and $\tau_?$.

Lemma 1. *For all P , $\vdash P : f \implies f(*)$ only contains τ and $\tau_?$ actions.*

Lemma 2. *$z \notin \text{fnv}(P) \wedge \Gamma \vdash P : f \implies f(z) = f(*)$*

A process type f takes a transition on a CO₂ prefix π when all its points $f(u)$ agree to take a transition on the abstract prefix $[\pi]_u$.

Definition 12. *We write $f \xrightarrow{\pi} f'$ whenever $\forall u \in \mathcal{N} \cup \mathcal{V} \cup \{*\}. f(u) \xrightarrow{[\pi]_u} f'(u)$.*

Example 9. Recall $P_1 = \text{do}_s \mathbf{a} + \text{do}_s \mathbf{b} + \text{do}_s \mathbf{z}$ from Ex. 2. Its typing is $\vdash P_1 : f = \lambda u . [\text{do}_s \mathbf{a}]_u + [\text{do}_s \mathbf{b}]_u + [\text{do}_s \mathbf{z}]_u$. Let $f' = \lambda u . \mathbf{0}$. We have that $f \xrightarrow{\text{do}_s \mathbf{a}} f'$, since:

- $[\text{do}_s \mathbf{a}]_s = \mathbf{a}$ and $f(s) = \mathbf{a} + \mathbf{b} + \mathbf{z} \xrightarrow{\mathbf{a}} \mathbf{0} = f'(s)$;
- $\forall v \neq s . [\text{do}_s \mathbf{a}]_v = \tau?$ and $f(v) = \tau? + \tau? + \tau? \xrightarrow{\tau?} \mathbf{0} = f'(v)$.

Note that, in this case, we also have $f \xrightarrow{\text{do}_s \mathbf{b}} f'$ and $f \xrightarrow{\text{do}_s \mathbf{z}} f'$.

If f is the type of some CO_2 process (i.e., f is *inhabited*), and a single point $f(u)$ takes an abstract transition, then the whole f can take a transition.

Lemma 3. $f(u) \xrightarrow{\alpha} T' \wedge f \text{ inhabited} \implies \exists \pi, f' . [\pi]_u = \alpha \wedge f'(u) = T' \wedge f \xrightarrow{\pi} f'$.

Definition 13 (Type honesty). f is honest iff $f(u)$ is honest, $\forall u \in \mathcal{N} \cup \mathcal{U} \setminus \{*\}$.

Note that, when $\vdash P : f$, checking the honesty of f amounts to checking $f(u)$ honest *only* for each $u \in \text{fv}(P)$. In fact, $f(v) = f(*)$ when $v \notin \text{fv}(P)$, and $f(*)$ is trivially honest because it cannot advertise contracts.

4.3 System typing

In this section we establish *type safety* for CO_2 processes: for any typeable closed P , $A[P]$ is honest. To prove this result, we have to consider the transitions of a process within a system. Hence, in order to construct an invariant of the system transitions (i.e., proving subject reduction), we extend typing to systems.

Type judgments for systems are of two kinds. A judgment of the form $\vdash_A S : f$ guarantees that a participant A in S behaves according to f . Instead, a judgment of the form $\vdash_A S \triangleright f$ means that A is *not* in S , and S is *compatible* with any participant A behaving as f . Our notion of compatibility is liberal: intuitively, it just checks that the context S does not contain forged contracts of A .

Definition 14 (System typing). The relations $\vdash_A S : f$ and $\vdash_A S \triangleright f$ are the smallest relations closed under the rules in Fig. 4.

Most rules in Fig. 4 are straightforward. Rules $[\text{T-SAFREE}^*]$ tell that A -free systems are compatible with all f . Rules $[\text{T-SFZ}^*]$ state that an f -compatible context (where f is the behaviour of A) may contain latent contracts of A if f realizes them. Rule $[\text{T-SFUSED}]$ is similar, but it deals with stipulated contracts of A . Rule $[\text{T-SDEL2}]$ is similar to $[\text{T-DEL}]$ for typing processes. Rule $[\text{T-SDEL1}]$ is dual: while in $[\text{T-SDEL2}]$ the type f abstracts the behaviour of A *within* S , in $[\text{T-SDEL1}]$ it represents the behaviour of A *outside* S . Note that if $A[P]$ is typeable, then it can be inserted in any A -free system, and the composed system will be typeable.

Example 10. Let $S_0 = B[Q] \mid C[\downarrow_x A \text{ says } c]$, with $B \neq A$ (note that S_0 is not A -free). Assume that $\vdash P : f$. Then, a typing derivation of $S = A[P] \mid S_0$ is:

$$\frac{\frac{\vdash P : f}{\vdash_A A[P] : f} \quad [\text{T-SA}]}{\frac{\frac{\frac{B \neq A}{\vdash_A B[Q] \triangleright f} \quad [\text{T-SAFREE1}] \quad \frac{f(x) \text{ realizes } c}{\vdash_A C[\downarrow_x A \text{ says } c] \triangleright f} \quad [\text{T-SFZ1}]}{\vdash_A B[Q] \mid C[\downarrow_x A \text{ says } c] = S_0 \triangleright f} \quad [\text{T-SPAR1}]}{\vdash_A A[P] \mid S_0 = S : f} \quad [\text{T-SPAR2}]}$$

Notice that S is typeable with f only if $f(x)$ realizes A 's contract c .

$$\begin{array}{c}
\frac{}{\vdash_A \mathbf{0} \triangleright f} \text{[T-SAFREE0]} \quad \frac{\mathbf{B} \neq \mathbf{A}}{\vdash_A \mathbf{B}[P] \triangleright f} \text{[T-SAFREE1]} \quad \frac{\mathbf{B} \neq \mathbf{A}}{\vdash_A \mathbf{C}[\downarrow_x \mathbf{B} \text{ says } c] \triangleright f} \text{[T-SAFREE2]} \quad \frac{\gamma \text{ A-free}}{\vdash_A s[\gamma] \triangleright f} \text{[T-SAFREE3]} \\
\frac{}{\vdash_A \mathbf{B}[\downarrow_s \mathbf{A} \text{ says } c] \triangleright f} \text{[T-SFzS]} \quad \frac{f(x) \text{ realizes } c}{\vdash_A \mathbf{B}[\downarrow_x \mathbf{A} \text{ says } c] \triangleright f} \text{[T-SFz1]} \quad \frac{\vdash_A \mathbf{B}[K] \triangleright f \quad \vdash_A \mathbf{B}[K'] \triangleright f}{\vdash_A \mathbf{B}[K \mid K'] \triangleright f} \text{[T-SFz2]} \\
\frac{\emptyset \vdash P: f}{\vdash_A \mathbf{A}[P]: f} \text{[T-SA]} \quad \frac{\vdash_A S \triangleright f\{f(*)/u\}}{\vdash_A (u)S \triangleright f} \text{[T-SDEL1]} \quad \frac{\vdash_A S: f \quad f(u) \text{ honest}}{\vdash_A (u)S: f\{f(*)/u\}} \text{[T-SDEL2]} \\
\frac{f(s) \text{ realizes } c}{\vdash_A s[\mathbf{A} \text{ says } c \mid \dots] \triangleright f} \text{[T-SFUSED]} \quad \frac{\vdash_A S \triangleright f \quad \vdash_A S' \triangleright f}{\vdash_A S \mid S' \triangleright f} \text{[T-SPAR1]} \quad \frac{\vdash_A S: f \quad \vdash_A S' \triangleright f}{\vdash_A S \mid S': f} \text{[T-SPAR2]}
\end{array}$$

Fig. 4. Typing rules for systems. Symmetric rules wrt \mid for [T-SFUSED], [T-SPAR2] omitted.

4.4 Type safety

Subject reduction guarantees that typeability is preserved by transitions. We need to distinguish two cases, according to which participant moves: either the participant \mathbf{A} under typing, or any other \mathbf{B} . If the transition is done by \mathbf{A} , then also its process type must take a transition, otherwise the type is preserved.

The substitution induced when executing a fuse also affects the type of the reduct system. For instance, consider $\mathbf{A}[P] \mid S$, where $\vdash P: f$ and $f(x) = T$. Assume that S fires a fuse, with a substitution σ mapping x to a fresh session name s . In the type, this is reflected by updating f according to σ , so to map s to T , and x , which is no longer free, to $f(*)$. Technically, this is done as follows.

Definition 15. For a type f and a mapping $\{v/\vec{u}\}$, the partial operator \bullet is:

$$f \bullet \{v/\vec{u}\} = \begin{cases} f & \text{if } \forall u_0 \in \vec{u}. f(u_0) = f(*) \\ f\{f(*)/u_0\}\{f(u_0)/v\} & \text{if } \exists! u_0 \in \vec{u}. f(u_0) \neq f(*) \end{cases}$$

Theorem 2 (Subject reduction). If $\vdash_A S: f$ with f honest, then:

$$\begin{aligned}
S \xrightarrow{\mathbf{A}: \pi, \sigma} S' &\implies \exists f'. f \xrightarrow{\pi} f' \wedge \vdash_A S': f' \bullet \sigma \\
S \xrightarrow{\mathbf{B}: \pi, \sigma} S' &\implies \vdash_A S': f \bullet \sigma \quad (\text{when } \mathbf{B} \neq \mathbf{A})
\end{aligned}$$

Progress guarantees that if a typeable process has a “non-blocking” type, then it can take a transition. More precisely, if the type of P on channel u can take a weak transition with label \mathbf{a} , then P will have \mathbf{a} in its weak ready do set.

Theorem 3 (Progress). For all $S \equiv s[\mathbf{A} \text{ says } c \mid \dots] \mid S'$, if $\vdash_A S: f$ with f honest, and $f(s) \xrightarrow{\mathbf{a}}_c^A$, then $\mathbf{a} \in \text{WRD}_s^A(S)$.

The main result of this paper is the type safety of CO_2 processes. It ensures that a participant \mathbf{A} with a well-typed process P will always respect her contracts — both those already advertised, and those that she will publish along her reductions. Therefore, \mathbf{A} will never be considered culpable in any context.

Theorem 4 (Type safety). $\forall \mathbf{A}[P]$ with P closed, $\vdash P: f \implies \mathbf{A}[P]$ is honest.

5 Concluding Remarks and Related Work

We have given a type system for a calculus of contracting processes. Type safety establishes *honesty*: typeable processes honour their contracts in all contexts.

In [3], A is considered “honest” when not definitely *culpable* (in any session), i.e., A eventually performs the actions her contract prescribes. In Def. 5, honesty is based on readiness, rather than culpability; we conjecture that the two notions are equivalent. The main advantage of this new approach compared to [3] is that it simplifies the proof of the correctness of the static analysis of honesty, by more directly relating abstract transitions with concrete ones. Also, the new definition helps in proving decidability of abstract honesty, which was left open in [3].

In [5] (multiparty) asserted global types are used to adapt design-by-contract to distributed interactions. In our framework, a participant declares its contract independently of the others; a CO₂ primitive (fuse) tries then to combine advertised contracts within a suitable agreement. In other words, one could think of our approach as based on orchestration rather than choreography.

In [16] the progress property is checked only when participants engage at most in one session at a time. Honesty is a sort of progress property, and our type system allows participants to interleave many sessions as done in [15]. A crucial difference with respect to [15] is that the typing discipline there requires the *consistency* of the local types of any two participants interacting in a session. Namely, if in a session s , A and B are typed as T_A and T_B respectively and they interact then the projection of T_A with respect to B must be dual of the projection of T_B with respect to A. In our type system instead, participants are typechecked ‘in isolation’ and to establish the honesty of a participant A our typing discipline only imposes that the surrounding context is A-free.

Other approaches deal with safety, by generating monitors that check at runtime the interactions of processes against their local contract (e.g., [13,12]).

The problem of checking if a contract c , representing the behaviour of a service, conforms to a role r of a choreography H has been investigated in [6]. Conformance of c is attained by establishing a *should testing* pre-order between c and the projection of H with respect to role r . Similar techniques have been used in [7] to define contract-based service composition. A main difference w.r.t. our approach, is that [6,7] do not consider conformance in the presence of dishonest participants. Actually, they focus on matching a contract as a role within a choreography, while we establish if a process abides by its own contracts, regardless of the context.

In [8,9], constraints are used to rule out “inconsistent” executions. This is orthogonal to our approach, since our aim is to blame participants that misbehave.

Acknowledgments. Work partially supported by Aut. Region of Sardinia under grants L.R.7/2007 CRP-17285 (TRICS), P.I.A. 2010 Project “Social Glue”, and by MIUR PRIN 2010-11 project “Security Horizons”, and by the Leverhulme Trust Programme Award “Tracing Networks”, and by COST Action IC1201: Behavioural Types for Reliable Large-Scale Software Systems (BETTY), and by European Union Seventh Framework Programme under grant agreement no. 295261 (MEALS).

References

1. M. Bartoletti, A. Scalas, E. Tuosto, and R. Zunino. Honesty by typing. Technical report, 2013. Available at <http://tcs.unica.it/publications>.
2. M. Bartoletti, E. Tuosto, and R. Zunino. Contract-oriented computing in CO₂. *Scientific Annals in Computer Science*, 22(1):5–60, 2012.
3. M. Bartoletti, E. Tuosto, and R. Zunino. On the realizability of contracts in dishonest systems. In *Proc. COORDINATION*, 2012.
4. M. Bartoletti and R. Zunino. A calculus of contracting processes. In *LICS*, 2010.
5. L. Bocchi, K. Honda, E. Tuosto, and N. Yoshida. A theory of design-by-contract for distributed multiparty interactions. In *CONCUR*, 2010.
6. M. Bravetti and G. Zavattaro. Towards a unifying theory for choreography conformance and contract compliance. In *Software Composition*, 2007.
7. M. Bravetti and G. Zavattaro. Contract-based discovery and composition of web services. In *Formal Methods for Web Services*. 2009.
8. M. Buscemi, M. Coppo, M. Dezani-Ciancaglini, and U. Montanari. Constraints for service contracts. In *Proc. TGC*, 2011.
9. M. G. Buscemi and U. Montanari. CC-Pi: A constraint-based language for specifying service level agreements. In *ESOP*, 2007.
10. S. Carpineti, G. Castagna, C. Laneve, and L. Padovani. A formal account of contracts for web services. In *WS-FM*, 2006.
11. G. Castagna, N. Gesbert, and L. Padovani. A theory of contracts for web services. *ACM TOPLAS*, 31(5), 2009.
12. T.-C. Chen, L. Bocchi, P.-M. Denilou, K. Honda, and N. Yoshida. Asynchronous distributed monitoring for multiparty session enforcement. In *Proc. TGC*. 2012.
13. T.-C. Chen and K. Honda. Specifying Stateful Asynchronous Properties for Distributed Programs. In *Proc. CONCUR*, 2012.
14. S. Christensen. *Decidability and Decomposition in Process Algebras*. PhD thesis, Edinburgh University, 1993.
15. M. Coppo, M. Dezani-Ciancaglini, N. Yoshida, and L. Padovani. Global Progress for Dynamically Interleaved Multiparty Sessions. 2013. Available at <http://www.di.unito.it/~padovani/Papers/CoppoDezaniYoshidaPadovani13.pdf>.
16. K. Honda, N. Yoshida, and M. Carbone. Multiparty asynchronous session types. In *POPL*, 2008.
17. R. Mayr. *Decidability and Complexity of Model Checking Problems for Infinite-State Systems*. PhD thesis, Technische Universität München, 1998.