

A Calculus of Contracting Processes

Massimo Bartoletti

Dipartimento di Matematica e Informatica, Università degli Studi di Cagliari

Roberto Zunino

Dipartimento di Ingegneria e Scienza dell'Informazione, Università degli Studi di Trento

Abstract

We propose a formal theory of contract-based computing. We model contracts as formulae in an intuitionistic logic extended with a “contractual” form of implication. Decidability holds for our logic: this allows us to mechanically infer the rights and the duties deriving from any set of contracts. We embed our logic in a core calculus of contracting processes, which combines features from concurrent constraints and calculi for multiparty sessions, while subsuming several idioms for concurrency.

1 Introduction

In Web transactions, the typical dynamics is that a client chooses a service provider that she trusts, relying on the fact that the service implements the required features. Such features are typically written in a “service level agreement” (SLA). Although this document is legally binding, it is not a formal specification. Formalising it would be desirable, for two main reasons. First, a formal SLA could be exploited by the client to mechanize the search of a service meeting her requests. Second, in the case the provider does not honour its SLA, automatic means could be devised to resolve the dispute. This would be more practical than taking legal steps against the provider, especially for transactions dealing with small amounts of money.

The interaction among parties has then to be regulated by a suitable *contract*, which formally subordinates the duties of a client to the duties of a service, and *vice versa*. The crucial problems are how to model a contract, how to infer when a set of contracts gives rise to an agreement among the stipulating parties, and how to single out the responsible of a possible violation.

An example. To give the intuition about our contracts, suppose there are two kids: Alice, who has a toy airplane, and Bob, who has a bike. Before sharing their toys, the two kids stipulate the following “gentlemen’s agreement”:

Alice: I will lend my airplane to you, Bob, provided that I borrow your bike.

Bob: I will lend my bike to you, Alice, provided that I borrow your airplane.

Let us write a for the atomic proposition “Alice lends her airplane” and b for “Bob lends his bike”. A (wrong) formalisation of the above contracts in classical logic could model Alice’s contract A as $b \rightarrow a$, and Bob’s B as $a \rightarrow b$. However, from this we cannot deduce the expected agreement, i.e. $A \wedge B \rightarrow a \wedge b$ does not hold. To solve this issue, we propose Propositional Contract Logic (PCL), that extends intuitionistic logic IPC with a *contractual implication* connective \twoheadrightarrow . In PCL we have the desired agreement:

$$(b \twoheadrightarrow a) \wedge (a \twoheadrightarrow b) \rightarrow a \wedge b$$

To put our contracts at work, we introduce a process calculus which embeds our logic. This calculus belongs to the family of concurrent constraints [35], using PCL formulae as constraints. A process can assert a constraint c (a PCL formula) through the primitive tell c . For instance, the following process models Alice exposing her contract:

$$(x) \text{ tell } b(x) \twoheadrightarrow a(x)$$

Formally, this will add $b(x) \twoheadrightarrow a(x)$ to the set of constraints. The formal parameter x represents the identifier of the actual session to be established between Alice and Bob. As it happens for sessions centered calculi [38, 14], sessions are an important aspect also in our calculus, since they allow for distinguishing among different instantiations of the same contract. The outer (x) is a scope delimitation for the variable x , similarly to the Fusion calculus [32].

After having exposed her contract, Alice will wait until finding that she has actually to lend her airplane to Bob. This is modelled as $\text{fuse}_x a(x)$. The primitive $\text{fuse}_x c$ implements a contract-based multiparty agreement. To do that, it checks the entailment of the constraint c , and binds the variable x to an actual session identifier, shared among all the

parties involved in the contract. So, we will model Alice as:

$$Alice = (x) (\text{tell } b(x) \multimap a(x). \text{fuse}_x a(x). \text{lendAirplane})$$

where the process *lendAirplane* (no further specified) models Alice actually lending her airplane to Bob. The overall behaviour of Alice is then: (i) issue the contract; (ii) wait until discovering the duty of lending the airplane; (iii) finally, lend the airplane. Dually, we model Bob as follows:

$$Bob = (y) (\text{tell } a(y) \multimap b(y). \text{fuse}_y b(y). \text{lendBike})$$

A possible interaction between Alice and Bob will be the following, where n stands for a fresh session identifier:

$$Alice|Bob \multimap^* (n) (\text{lendAirplane}\{n/x\} | \text{lendBike}\{n/y\})$$

As expected, the resulting process shows Alice and Bob actually sharing their toys, in the session identified by n .

The logic PCL also allows for a more precise model of the above scenario, by linking the contracts with the identity of the principals issuing them (see Ex. 3). This information can be exploited in our calculus to automatically detect the responsible of a violation (see e.g. Ex. 6).

Contributions. We propose the logic PCL, which extends IPC with contractual implication. We provide it with an Hilbert-style axiomatisation and a Gentzen-style sequent calculus, which we prove equivalent. We study the relations between PCL, IPC and classical logic. The main results about PCL are cut elimination and decidability. We implement a proof search algorithm for PCL [37], also including an extension with a lax modality, to explicitly link contracts and principals. We then exploit PCL as a basic building block for designing a calculus of contracting processes. Our calculus is expressive enough to encode several concurrency idioms, among which Linda, the π -calculus and graph rewriting. We show our logic and calculus applicable to model real-world scenarios through several examples.

Because of space constraints, we include all the proofs, the encodings, as well as further results and examples about our logic and calculus, in two Technical Reports [6, 5].

2 A Logic for Contracts

Desirable properties. We start by characterizing our logic through a set of properties that we would expect to be enjoyed by any logic for contracts.

As shown in Sect. 1, a basic property of contractual implication is that of allowing two dual contracting parties to “handshake”, so to make their agreement effective. This is resumed by the following *handshaking* property:

$$\vdash (p \multimap q) \wedge (q \multimap p) \multimap p \wedge q \quad (1)$$

A generalisation of the above property to the case of n contracting parties is also desirable. It is a sort of “circular” handshaking, where the $(i + 1)$ -th party, in order to promise some duty p_{i+1} , relies on a promise p_i made by the i -th party (in a circular fashion, the first party relies on the promise of the last one). In the case of n parties, we expect:

$$\vdash (p_1 \multimap p_2) \wedge \cdots \wedge (p_{n-1} \multimap p_n) \wedge (p_n \multimap p_1) \multimap p_1 \wedge \cdots \wedge p_n \quad (2)$$

As a concrete example, consider an e-commerce scenario where a client C can buy items from a seller S , and pay them through a credit card. To mediate the interaction between C and S , there is a bank B which manages payments. The contracts issued by the three parties could be as follows:

Client: I will click “pay” provided that my item is shipped

Seller: I will ship your item provided that I get the money

Bank: I will transfer money to the seller provided that the client clicks “pay”.

Let the atomic propositions ship, click, and pay denote, respectively, the facts “seller ships item”, “client clicks pay”, and “bank transfers money”. The above contracts can then be modelled as:

$$C = \text{ship} \multimap \text{click} \quad B = \text{click} \multimap \text{pay} \quad S = \text{pay} \multimap \text{ship}$$

Then, by property (2) we deduce a successful transaction:

$$\vdash C \wedge B \wedge S \multimap \text{pay} \wedge \text{ship}$$

Note that, in the special case $n = 1$, the above “circular” handshaking property turns into a particularly simple form:

$$\vdash (p \multimap p) \multimap p \quad (3)$$

Intuitively, (3) models that promising p provided that p , implies p (actually, also the converse holds, so that promise is equivalent to p). It also follows from (1) when $p = q$.

Another generalisation of the toy-exchange scenario to the case of n kids is also desirable. It is a sort of “greedy” handshaking, because now a party promises p_i only provided that *all* the other parties promise their duties, i.e. $p_1, \dots, p_{i-1}, p_{i+1}, \dots, p_n$.

$$\vdash \bigwedge_{i \in 1..n} ((p_1 \wedge \cdots \wedge p_{i-1} \wedge p_{i+1} \wedge \cdots \wedge p_n) \multimap p_i) \multimap p_1 \wedge \cdots \wedge p_n \quad (4)$$

We will now focus on further logical properties of contractual implication. As shown by (1), a contract $p \multimap q$ becomes effective, i.e. implies the promise q , when matched by a dual contract $q \multimap p$. Even more directly, $p \multimap q$ should be effective also when the premise p is already true:

$$\vdash p \wedge (p \multimap q) \multimap q \quad (5)$$

In other words, contractual implication should be *stronger* than standard implication, i.e.:

$$\vdash (p \multimap q) \rightarrow (p \rightarrow q) \quad (6)$$

On the other hand, we do not want that also the converse holds, since this would equate the two forms of implication: that is, $\not\vdash (p \rightarrow q) \rightarrow (p \multimap q)$.

We want contractual implication to share with standard implication a number of properties. First, a contract that promises true (written \top) is always satisfied, regardless of the precondition. So, we expect the following tautology:

$$\vdash p \multimap \top \quad (7)$$

Differently from standard implication, we do not want a contract with a false precondition (written \perp) to always hold, i.e. $\not\vdash \perp \rightarrow p$. To see why, assume $\perp \rightarrow p$ is a tautology, for all p . Then, it would also be the case for $p = \perp$, so (3) would deduce a contradiction: $(\perp \rightarrow \perp) \rightarrow \perp$.

We want \rightarrow to enjoy transitivity, similarly to \multimap :

$$\vdash (p \rightarrow q) \wedge (q \rightarrow r) \rightarrow (p \rightarrow r) \quad (8)$$

Back to our previous example, transitivity would allow the promise of the client (ship \rightarrow click) and that of the bank (click \rightarrow pay) to be combined in the promise ship \rightarrow pay.

Contractual implication should also enjoy a stronger form of transitivity. We illustrate it with the help of an example. Suppose an air-flight customer who wants to book a flight. The customer contract promises to pay the required amount, provided that she obtains a flight reservation. Suppose now that an airline company starts a special offer, in the form of a free drink for each customer.

$$\begin{aligned} \text{Customer} &= \text{bookFlight} \multimap \text{pay} \\ \text{AirLine} &= \text{pay} \multimap \text{bookFlight} \wedge \text{freeDrink} \end{aligned}$$

Of course, the two contracts should give rise to an agreement, because the airline company is promising a better service than the one required by the customer contract. We then expect to be able to “weaken” the *AirLine* contract:

$$\vdash \text{AirLine} \rightarrow (\text{pay} \multimap \text{bookFlight})$$

Alternatively, one could make the two contracts match by making stronger the precondition required by the customer:

$$\vdash \text{Customer} \rightarrow (\text{bookFlight} \wedge \text{freeDrink} \multimap \text{pay})$$

More in general, we want the following two properties hold for any logic for contracts. They say that the promise in a contract can be arbitrarily weakened (9), while the precondition can be arbitrarily strengthened (10).

$$\vdash (p \multimap q) \wedge (q \rightarrow q') \rightarrow (p \multimap q') \quad (9)$$

$$\vdash (p' \rightarrow p) \wedge (p \multimap q) \rightarrow (p' \multimap q) \quad (10)$$

Note that (8), (9), (10) cover three of the four possible cases of transitivity which mix standard and contractual implications. The fourth case would, instead, make the two forms of implication coincide, so it is *not* a desirable property.

Another desirable property is that, if a promise q is already true, then it is also true any contract which promises q :

$$\vdash q \rightarrow (p \multimap q) \quad (11)$$

Of course, we do not want the converse to hold: a contract not always implies its promise: $\not\vdash (p \multimap q) \rightarrow q$.

Syntax. The syntax of PCL extends that of IPC. It includes the standard connectives $\neg, \wedge, \vee, \rightarrow$ and the contractual implication \multimap . We assume a denumerable set $\{p, q, r, s, \dots\}$ of prime (atomic) formulae. PCL formulae are denoted with the letters p, q, r, s, \dots (note that the font differs from that used for prime formulae). The precedence of IPC operators is, from highest to lowest: $\neg, \wedge, \vee, \rightarrow$. We stipulate that \rightarrow has the same precedence as \multimap .

Definition 1 *The formulae of PCL are defined as:*

$$p ::= \perp \mid \top \mid p \mid \neg p \mid p \vee p \mid p \wedge p \mid p \rightarrow p \mid p \multimap p$$

We let $p \leftrightarrow q$ be syntactic sugar for $(p \rightarrow q) \wedge (q \rightarrow p)$.

We now present an Hilbert-style axiomatization for PCL.

Definition 2 *The proof system of PCL comprises all the axioms of IPC, the Modus Ponens rule CUT, and the axioms:*

$$\begin{aligned} \top \multimap \top & \quad \text{[ZERO]} \\ (p \multimap p) \rightarrow p & \quad \text{[FIX]} \\ (p' \rightarrow p) \rightarrow (p \multimap q) \rightarrow (q \rightarrow q') \rightarrow (p' \multimap q') & \quad \text{[PREPOST]} \end{aligned}$$

The above axioms are a subset of the properties discussed above. The axiom ZERO is a subcase of (7), FIX is just (3), while the axiom PREPOST combines (9) and (10). As expected, this set of axioms is actually sound and complete w.r.t. all the properties marked above as desirable.

Lemma 1 *Properties (1-11) are theorems of PCL. Also:*

$$\begin{aligned} & \vdash (p \multimap q) \wedge (q \multimap r) \rightarrow (p \multimap (q \wedge r)) \\ & \vdash (p \multimap (q \wedge r)) \rightarrow (p \multimap q) \wedge (p \multimap r) \\ & \vdash (p \multimap q) \vee (p \multimap r) \rightarrow (p \multimap (q \vee r)) \\ & \vdash (p \multimap q) \rightarrow ((q \rightarrow p) \rightarrow q) \end{aligned}$$

We present below some of the most significant results about our logic. For a more comprehensive account, including detailed proofs of all our results, see [6].

First, PCL is consistent. Also, negation-free formulae do not lead to inconsistencies.

Theorem 1 PCL is consistent, i.e. $\not\vdash \perp$. Also, if p is free from $\{\perp, \neg\}$, then $\not\vdash p \rightarrow \perp$.

As expected, the following are *not* tautologies of PCL :

$$\begin{array}{ll} \not\vdash (p \rightarrow q) \rightarrow (p \multimap q) & \not\vdash (p \multimap q) \rightarrow q \\ \not\vdash \perp \multimap p & \not\vdash ((q \rightarrow p) \rightarrow q) \rightarrow (p \multimap q) \end{array}$$

Note that if we augment our logic with the axiom of excluded middle, then $(p \multimap q) \leftrightarrow q$ becomes a theorem. This would make contractual implication coincide with right projection, so losing most of the intuition behind mutual agreements. For this reason we use IPC, instead of classical logic, as the basis of PCL .

Another main result about PCL is its decidability. To prove that, we have devised a Gentzen-style sequent calculus, equivalent to the Hilbert-style axiomatisation. In particular, we have extended the sequent calculus for IPC presented in [33] with rules for the contractual implication \multimap . Note that PREPOST introduces \multimap on the right, and eliminates it on the left (similarly e.g. to $\circ L$ of lax logic [19]).

Definition 3 The sequent calculus of PCL includes all the rules for IPC [6], and the following additional rules.

$$\begin{array}{ccc} \text{[ZERO]} & \text{[FIX]} & \text{[PREPOST]} \\ \frac{\Gamma \vdash q}{\Gamma \vdash p \multimap q} & \frac{\Gamma, p \multimap q, r \vdash p \quad \Gamma, p \multimap q, q \vdash r}{\Gamma, p \multimap q \vdash r} & \frac{\Gamma, p \multimap q, a \vdash p \quad \Gamma, p \multimap q, q \vdash b}{\Gamma, p \multimap q \vdash a \multimap b} \end{array}$$

We now establish the equivalence between the two logical systems of PCL. In the following theorem, we denote with \vdash_H provability in the Hilbert-style system, while \vdash_G is used for Gentzen-style provability.

Theorem 2 For all PCL formulae p : $\vdash_H p \iff \emptyset \vdash_G p$.

Our sequent calculus enjoys cut elimination. The proof is non-trivial, since the rules for \multimap are not dual, unlike e.g. left/right rules for \wedge . Nevertheless, the structural approach of [33] can be adapted. A cut on a formula p is replaced by cuts on strict subformulae of p , and cuts on p having a shorter proof tree. Some insightful cases of our proof, as well as the induction metric used, are in Appendix A; the full details are in [6].

Theorem 3 (Cut Elimination) If p is provable in PCL, then there exists a proof of p not using the CUT rule.

The subformula property holds in PCL. Cut-free proofs only involve subformulae of the sequent at hand.

Theorem 4 (Subformula Property) Let D be a cut-free proof of $\Gamma \vdash p$. Then, the formulae occurring in D are subformulae of those occurring in Γ and p .

Theorems 3 and 4 allow for exhaustively searching the proof space, so implying decidability.

Theorem 5 The logic PCL is decidable.

As a further support to our logic, we have implemented a proof search algorithm [37], which decides whether any given formula is a tautology or not. Despite the problem being PSPACE complete [36], the performance of our tool is acceptable for the examples presented in this paper.

We now establish some expressiveness results, relating PCL and IPC. More in detail, we consider whether sound and complete homomorphic encodings exist, i.e. whether \multimap can be regarded as syntactic sugar for some IPC context.

Definition 4 A homomorphic encoding m is a function from PCL formulae to IPC formulae such that: m is the identity on prime formulas, \top , and \perp ; it acts homomorphically on $\wedge, \vee, \rightarrow, \neg$; it satisfies $m(p \multimap q) = \mathcal{C}[m(p), m(q)]$ for some fixed IPC context $\mathcal{C}(\bullet, \bullet)$.

Of course, each homomorphic encoding is uniquely determined by the context \mathcal{C} . Several *complete* encodings exist:

Lemma 2 The following homomorphic encodings are complete, i.e. they satisfy $\vdash p \implies \vdash_{IPC} m_i(p)$. Also, they are pairwise non-equivalent in IPC (the prime a is arbitrary).

$$\begin{array}{l} m_0(p \multimap q) = m_0(q) \\ m_1(p \multimap q) = (m_1(q) \rightarrow m_1(p)) \rightarrow m_1(q) \\ m_2(p \multimap q) = \neg\neg(m_2(q) \rightarrow m_2(p)) \rightarrow m_2(q) \\ m_3(p \multimap q) = \neg(m_3(q) \rightarrow m_3(p)) \vee m_3(q) \\ m_4(p \multimap q) = ((m_4(q) \rightarrow m_4(p)) \vee a) \rightarrow m_4(q) \end{array}$$

However, there can be no *sound* encodings, so \multimap is not just syntactic sugar. Indeed, a sound encoding would allow us to derive Peirce's axiom in PCL, violating the fact that PCL conservatively extends IPC [6].

Theorem 6 If m is a homomorphic encoding of PCL into IPC, then m is not sound, i.e. there exists a PCL formula p such that $\vdash_{IPC} m(p)$ and $\not\vdash p$.

In [6] we have proved further properties of PCL, including some relations between PCL and IPC, the modal logic S4, and propositional lax logic. Also, we have explored there further application scenarios for our logic.

Example 1 (Online sale) We describe a possible online sale between two parties. To buy an item, the buyer contacts the bank, to reserve from his account a given amount of money for the transaction. When this happens (modelled with the prime formula *lock*), that amount is no longer available. Then, the buyer makes an offer to the seller (offer). When provided with a good offer, and the money has been reserved, the seller will send the item (send). Otherwise, she cancels the transaction (abort). When the transaction is aborted, the bank cancels the reservation (unlock), so the money can be reused.

Formally, the buyer agrees to $\text{lock} \wedge \text{offer}$, provided that either the item is sent, or the reservation is cancelled. The seller agrees to evaluate the offer. The bank agrees to cancel the reservation when the transaction is aborted.

$$\begin{aligned} \text{Buyer} &= (\text{send} \vee \text{unlock}) \rightarrow (\text{lock} \wedge \text{offer}) \\ \text{Seller} &= \text{offer} \rightarrow ((\text{lock} \rightarrow \text{send}) \vee \text{abort}) \\ \text{Bank} &= (\text{lock} \wedge \text{abort}) \rightarrow \text{unlock} \end{aligned}$$

Under these assumptions, either the item is sent, or the transaction is aborted and the reservation cancelled:

$$\vdash (\text{Buyer} \wedge \text{Seller} \wedge \text{Bank}) \rightarrow (\text{send} \vee (\text{abort} \wedge \text{unlock}))$$

Example 2 (Dining retailers) Around a table, n cutlery retailers are about to have dinner. At the center of the table, there is a large dish of food. Despite the food being delicious, the retailers cannot start eating right now. To do that, and follow the proper etiquette, each retailer needs a complete cutlery set, consisting of n pieces of different kinds. Each of the n retailers owns a distinct set of n piece of cutlery, all of the same kind. The retailers start discussing about trading their cutlery, so that they can finally eat.

We formalize this scenario as follows. We number the retailers r_1, \dots, r_n together with the kinds of pieces of cutlery, so that r_i initially owns n pieces of kind i . We then write $g_{i,j}$ for “ r_i gives a piece (of kind i) to r_j ”. Since retailers can use their own cutlery, we assume $g_{i,i}$ to be true. Retailer r_i can start eating whenever $e_i = \bigwedge_j g_{j,i}$. Instead, he provides the cutlery to others whenever $p_i = \bigwedge_j g_{i,j}$.

Suppose that r_1 commits to a simple exchange with r_2 : they commit to $g_{2,1} \rightarrow g_{1,2}$ and $g_{1,2} \rightarrow g_{2,1}$, and the exchange takes place since $g_{2,1} \wedge g_{1,2}$ can be derived. While this seems a fair deal, it actually exposes r_1 to a risk: if r_3, \dots, r_n perform a similar exchange with r_2 , then $g_{2,i} \wedge g_{i,2}$ for all i . In particular, $g_{i,2}$ holds for all i , so r_2 can start eating. This is however not necessarily the case for r_1 , since r_3 has not committed to any exchange with r_1 .

A wise retailer would then never agree to a simple exchange $g_{2,1} \rightarrow g_{1,2}$. Instead, the retailer r_1 could commit to the following safer contract:

$$e_1 \rightarrow p_1 = g_{1,1} \wedge g_{2,1} \wedge \dots \wedge g_{n,1} \rightarrow g_{1,1} \wedge g_{1,2} \wedge \dots \wedge g_{1,n}$$

The idea is simple: r_1 requires each piece of cutlery, i.e. r_1 requires to be able to start eating (e_1). When this happens, r_1 agrees to provide each other retailer with a piece of his cutlery (p_1). If each retailer r_i commits to the analogous contract, we have the desired agreement (proof in [6]).

$$\vdash \bigwedge_i (e_i \rightarrow p_i) \rightarrow \bigwedge_i e_i$$

Principals. As illustrated by the examples above, PCL allows for inferring whether some promise is implied by a set of contracts. In real-world scenarios, each contract will be

issued by a given *principal*. It would then be useful to represent the binding between principals and contracts within the logic. This would allow, for instance, to single out the principal who is responsible for a violation, and possibly to take countermeasures against him. To this aim, we extend our logic with a *says* modality, similarly to [21].

Definition 5 The syntax of PCL^{says} extends that of PCL (Def. 1) with the construct $a \text{ says } p$, where we assume a set of (atomic) principals, ranged over by a, b, \dots

The formula $a \text{ says } p$ represents the fact that the principal a has issued a contract p . We now develop the proof theory of PCL^{says} . Essentially, we extend the PCL axioms with those of the logic ICL [21]. This is an indexed lax logic, where the lax modality corresponds to *says*. Remarkably, this extension preserves all the main results of PCL, in particular its decidability.

Definition 6 The Hilbert-style axiomatisation of PCL^{says} extends that of PCL (Def. 2) with the following axioms:

$$\begin{aligned} p \rightarrow (a \text{ says } p) & \quad [\text{SAYSR}] \\ (a \text{ says } a \text{ says } p) \rightarrow a \text{ says } p & \quad [\text{SAYSM}] \\ (p \rightarrow q) \rightarrow (a \text{ says } p) \rightarrow (a \text{ says } q) & \quad [\text{SAYSF}] \end{aligned}$$

Example 3 We now make explicit the binding between the duties and the principals of the toy exchange example of Sect. 1. Alice commits herself to lend her airplane, provided that Bob commits himself to lend his bike (and vice versa).

$$\begin{aligned} p_{\text{toy}} &= \text{Alice says } ((\text{Bob says } b) \rightarrow a) \wedge \\ &\quad \text{Bob says } ((\text{Alice says } a) \rightarrow b) \end{aligned}$$

Such contract implies the expected duties:

$$\vdash p_{\text{toy}} \rightarrow \text{Alice says } a \wedge \text{Bob says } b$$

The above duties can be exploited by a third party (a sort of “automated” judge) which has to investigate the responsibilities of the involved parties, in the unfortunate case that the contract is not respected. For instance, if our judge is given the evidence that Alice’s airplane has never been lent to Bob, then he will infer that Alice has not respected her contract (and possibly punish her), that is:

$$p_{\text{toy}} \wedge \neg a \rightarrow (\text{Alice says } a) \wedge \neg a \rightarrow \text{Alice says } \perp$$

In Ex. 6 below, we shall see how the notion of principals is exploited in our calculus of contracting processes.

Explicitly representing principals has some additional benefits, especially when putting our logic at work in insecure environments populated by attackers. Actually, an attacker could maliciously issue a “fake” contract, where he

makes a promise that he cannot actually implement, e.g. because the promised task can only be performed by another party. By binding each contract with its principal, it is easy to realize when someone has attempted such a fraud, because the principal who has signed the contract is different from who is due to implement the promised behaviour.

Back to our technical development, all the main results of PCL are also enjoyed by $\text{PCL}^{s\text{ays}}$ (see [6] for details).

Theorem 7 *If p is provable in $\text{PCL}^{s\text{ays}}$, then there exists a proof of p without CUT. Also, $\text{PCL}^{s\text{ays}}$ is decidable.*

3 A contract calculus

We now define our calculus. We use a denumerable set of names n, m, \dots and a denumerable set of variables x, y, \dots . Metavariables a, b range over both names and variables. Intuitively, names play the same role as in the π -calculus, while variables roughly behave as names in the fusion calculus [32]. Distinct names represent distinct concrete objects, each one with its own identity. Instead, distinct variables can be *fused*, by instantiating them to the same name. Unlike [32], our calculus can fuse a variable only once.

Syntax. We extend prime formulae of PCL with names and variables as parameters: $p(a, b)$. Note that we do not introduce quantifiers in PCL, which then remains a propositional logic. Indeed, the formula $p(a, b)$ is still atomic from the point of view of the logic. Similarly, names and variables can appear as principals in a *says* p . Intuitively, here names model known principals, while variables still-unknown ones (see e.g. Ex. 6). We let letters c, d range over formulae, while letters u, v range over $\{\perp, \neg\}$ -free formulae. The syntax of our contract calculus follows.

$$\begin{aligned} \pi &::= \tau \mid \text{tell } u \mid \text{check } c \mid \text{ask}_{\vec{x}} c \mid \text{fuse}_x c & (\text{prefixes}) \\ P &::= u \mid \sum_{i \in I} \pi_i.P_i \mid P|P \mid (a)P \mid X(\vec{a}) & (\text{processes}) \end{aligned}$$

Processes are mostly standard, and include active constraints u , sum of guarded processes $\sum_i \pi_i.P_i$, parallel composition $P|P$, scope delimitation $(a)P$. We use a set of definitions $\{X_i(\vec{x}) \doteq P_i\}_i$ with the provision that each occurrence of X_j in P_k is guarded, i.e. behind some prefix.

We write 0 for the empty sum, and we use $+$ to merge sums, that is: $\sum_{i \in I} \pi_i.P_i + \sum_{i \in J} \pi_i.P_i = \sum_{i \in I \cup J} \pi_i.P_i$ when $I \cap J = \emptyset$. Singleton sums are simply written $\pi.P$.

The prefix τ is the standard silent operation of CCS. The prefixes *tell* and *check* are those of Concurrent Constraints (CC) [35]. A *tell* u augments the context with the (negation-free) formula u . By Theorem 1, the context will always be consistent. A *check* c checks if c is consistent with the context. The prefix $\text{ask}_{\vec{x}} c$ generalizes the prefix $\text{ask } c$ from CC: actually, they are equivalent when $\vec{x} = \emptyset$. An $\text{ask}_{\vec{x}} c$ stops a process until formula c can be deduced from the context.

To make that happen, variables \vec{x} are suitably instantiated to names. The prefix $\text{fuse}_x c$ is peculiar of our contract calculus, and drives the fusion of variables. Like *ask*, it stops a process until instantiating some variables to names makes c deducible. Note that the set of such variables is not specified in the prefix (as it was for $\text{ask}_{\vec{x}}$); instead, it is inferred from the context according to the *local minimal fusion* policy, to be introduced in Def. 7. The variable x in $\text{fuse}_x c$ is instantiated to a fresh name, that intuitively represents a fresh session identifier. Indeed, the intended use of a $\text{fuse}_x c$ is to initiate a new session, by accepting a contract which implies c . To do that, x is replaced by a fresh session ID, and some variables in the context are possibly instantiated to names (e.g. to bind unknown principals to actual ones). Instead, an *ask* can be used to join an already initiated session (no fresh ID is then generated).

Free variables and names, as usual, are those not under a delimitation. Alpha conversion and substitutions are defined accordingly. As a special case, when a variable x is instantiated to a name, the prefix $\text{fuse}_x c$ behaves as a plain $\text{ask } c$, i.e. $(\text{fuse}_x c)\{n/x\} = \text{ask}(c\{n/x\})$. Also, $(\text{ask}_{\vec{x}\vec{y}} c)\{n/x\} = \text{ask}_{\vec{y}}(c\{n/x\})$. Henceforth, we consider processes up-to alpha conversion.

Semantics. We now define the semantics of processes through a two-layered transition system (Fig. 3). The bottom layer is an LTS $\xrightarrow{\alpha}$ between processes, which provides a compositional semantics. Actions α are as follows, where C denotes a set of PCL formulae.

$$\alpha ::= \tau \mid C \mid C \vdash_{\vec{x}}^A c \mid C \vdash_x^F c \mid C \not\vdash \perp \mid (a)\alpha \quad (\text{actions})$$

The action τ represents an internal move. The action C is an advertisement of a set of active constraints. The action $C \vdash_{\vec{x}}^A c$ is a *tentative* action, generated by a process attempting to fire an $\text{ask}_{\vec{x}} c$ prefix. This action also carries a set C as the collection of active constraints discovered so far. Similarly for $C \vdash_x^F c$ and $\text{fuse}_x c$, as well as for $C \not\vdash \perp$ and $\text{check } c$. In the last case, C also includes c . The delimitation in $(a)\alpha$ is for scope extrusion, as in the labelled semantics of the π -calculus. We write $(\vec{a})\alpha$ for a set of distinct delimitations, neglecting their order, e.g. $(ab) = (ba)$. We simply write $(\vec{a}\vec{b})$ for $(\vec{a} \cup \vec{b})$.

The first two lines of Fig. 3 handle the base cases of our semantics. The rule **TAU** allows a τ prefix to fire. The rules **ASK**, **FUSE** and **CHECK** simply generate the corresponding actions. The rule **TELL** adds a constraint to the environment, thus making it *active*. Active constraints can then signal their presence through the **CONSTR** rule: each constraint generates its own singleton. A sum of guarded processes $\sum \pi.P$ can instead signal that it is *not* a constraint, by generating the empty set of constraints through **IDLESUM**.

The rules **SUM**, **DEF** (in the third line) are quite standard: they handle external choice and (possibly recursive) defini-

$$\begin{array}{c}
\tau.P \xrightarrow{\tau} P \text{ [TAU]} \quad \text{ask}_{\vec{x}} c.P \xrightarrow{\emptyset \vdash_{\vec{x}}^A c} P \text{ [ASK]} \quad \text{fuse}_{x,c}.P \xrightarrow{\emptyset \vdash_x^F c} P \text{ [FUSE]} \quad \text{check } c.P \xrightarrow{\{c\} \not\vdash \perp} P \text{ [CHECK]} \\
\text{tell } u.P \xrightarrow{\tau} u|P \text{ [TELL]} \quad u \xrightarrow{\{u\}} u \text{ [CONSTR]} \quad \sum_i \pi_i.P_i \xrightarrow{\emptyset} \sum_i \pi_i.P_i \text{ [IDLESUM]} \\
\frac{\pi_j.P_j \xrightarrow{\alpha} P'}{\sum_i \pi_i.P_i \xrightarrow{\alpha} P'} \text{ [SUM]} \quad \frac{P\{\vec{a}/\vec{x}\} \xrightarrow{\alpha} P'}{X(\vec{a}) \xrightarrow{\alpha} P'} \text{ if } X(\vec{x}) \doteq P \text{ [DEF]} \quad \frac{P \xrightarrow{\alpha} P'}{(a)P \xrightarrow{\alpha} (a)P'} \text{ if } a \notin \alpha \text{ [DEL]} \quad \frac{P \xrightarrow{\alpha} P'}{(a)P \xrightarrow{(a)\alpha} P'} \text{ [OPEN]} \\
\frac{P \xrightarrow{(\vec{a})C} P' \quad Q \xrightarrow{(\vec{b})C'} Q'}{P|Q \xrightarrow{(\vec{a}\vec{b})(C \cup C')} P'|Q'} \dagger \text{ [PARCONSTR]} \quad \frac{P \xrightarrow{(\vec{a})C} P' \quad Q \xrightarrow{(\vec{b})(C' \not\vdash \perp)} Q'}{P|Q \xrightarrow{(\vec{a}\vec{b})(C \cup C' \not\vdash \perp)} P'|Q'} \dagger \text{ [PARCHECK]} \quad \frac{P \xrightarrow{\tau} P'}{P|Q' \xrightarrow{\tau} P'|Q'} \dagger \text{ [PARTAU]} \\
\frac{P \xrightarrow{(\vec{a})C} P' \quad Q \xrightarrow{(\vec{b})(C' \vdash_{\vec{x}}^A c)} Q'}{P|Q \xrightarrow{(\vec{a}\vec{b})(C \cup C' \vdash_{\vec{x}}^A c)} P'|Q'} \dagger \text{ [PARASK]} \quad \frac{P \xrightarrow{(\vec{a})C} P' \quad Q \xrightarrow{(\vec{b})(C' \vdash_x^F c)} Q'}{P|Q \xrightarrow{(\vec{a}\vec{b})(C \cup C' \vdash_x^F c)} P'|Q'} \dagger \text{ [PARFUSE]} \quad (\dagger) \begin{array}{l} \vec{a} \text{ fresh in } \vec{b}, C', c, x, \vec{x}, Q' \\ \vec{b} \text{ fresh in } C, P' \end{array} \\
\frac{P \xrightarrow{(\vec{x}\vec{n}\vec{a})(C \vdash_{\vec{x}}^A c)} P'}{P \xrightarrow{\tau} (\vec{n}\vec{a})P'\sigma} \text{ if } \begin{array}{l} C\sigma \vdash c\sigma \\ \sigma(\vec{x}) \subseteq \vec{n} \end{array} \text{ [CLOSEASK]} \quad \frac{P \xrightarrow{(x\vec{y}\vec{m}\vec{a})(C \vdash_x^F c)} P'}{P \xrightarrow{\tau} (n\vec{m}\vec{a})P'\sigma} \begin{array}{l} C \vdash^\sigma c \\ \sigma(x) = n \text{ fresh} \\ \sigma(\vec{y}) \subseteq n\vec{m} \end{array} \text{ [CLOSEFUSE]} \\
\frac{P \xrightarrow{\tau} P'}{P \rightsquigarrow P'} \text{ [TOPTAU]} \quad \frac{P \xrightarrow{(\vec{a})(C \not\vdash \perp)} P'}{P \rightsquigarrow (\vec{a})P'} \text{ if } C \not\vdash \perp \text{ [TOPCHECK]}
\end{array}$$

Figure 1. The transition system for the contract calculus (symmetric rules for $|$ are omitted).

tions, respectively. The rules DEL, OPEN handle delimitation. As usual, when a is not mentioned in an action, we can propagate the action across a delimitation (a) using DEL. The rule OPEN instead allows for scope extrusion. Note that OPEN has no side conditions: the checks needed for scope extrusion are handled by the PAR* rules.

The next two lines of rules handle parallel composition. The rule PARCONSTR merges the sets of constraints advertised by two parallel processes. The rule PARASK allows for augmenting the constraints C in the tentative action $C \vdash_{\vec{x}}^A c$ generated by an ASK, by also accounting for the set of constraints advertised by the parallel process. Similarly for the rules PARFUSE and PARCHECK. The rule PARTAU simply propagates τ actions of parallel processes. The PAR* rules also merge the set of delimitations; variable and name captures are avoided through the side condition (\dagger).

The CLOSE* rules are the crucial ones, since they provide the mechanism to finalize the actions generated by ask and fuse. The rule CLOSEASK instantiates the variables \vec{x} (originally in a prefix $\text{ask}_{\vec{x}} c$) to a subset of the names \vec{n} collected so far. This is done through a substitution σ . When $c\sigma$ is deducible under the constraint set $C\sigma$, a silent action τ is generated, the delimitation $(\vec{n}\vec{a})$ is brought back to the process level, and σ is applied to the residual process P' . Note that the restriction (\vec{x}) is no longer needed, as all the variables in \vec{x} have been instantiated by σ . The rule CLOSEFUSE instantiates a subset \vec{y} of the variables collected in the action to a set of names. These names comprise a subset \vec{m} of the collected ones, plus a *fresh* name n . The variable x (originally in a prefix $\text{fuse}_{x,c}$) is instantiated to n . A transition is then possible when the substitution σ , defining the instantiation, makes the formula c deducible from the constraint set C under the *local minimal fusion* relation \vdash^σ , to be de-

finied in a while (Def. 7). When this happens, a silent action τ is generated, the delimitation $(\vec{n}\vec{m}\vec{a})$ is brought back to the process level, and σ is applied to the residual process P' . Apart from generating a fresh name, another key difference w.r.t. CLOSEASK is that the set of variables to be instantiated is chosen rather arbitrarily by CLOSEFUSE, while it is hard-wired in a prefix in CLOSEASK. This is the reason why CLOSEFUSE checks local minimal fusion, and not just deducibility, to further restrict the choice of σ .

Summing up, the LTS given by $\xrightarrow{\alpha}$ allows for the generation of tentative actions for ask and fuse (rules ASK, FUSE), which can then be converted to $\xrightarrow{\tau}$ whenever enough constraints are discovered (rules CLOSE*). Then, the τ action can be propagated towards the top level (rule PARTAU). A prefix check c , instead, cannot be handled in the same fashion, since it requires to check the consistency of c with respect to *all* the active constraints. To this aim, we use the reduction relation \rightsquigarrow , layered over the $\xrightarrow{\alpha}$ relation. The reduction \rightsquigarrow only includes internal moves $\xrightarrow{\tau}$ (rule TOPTAU) and successful check moves (TOPCHECK). This effectively discards tentative actions, filtering the unsuccessful ones.

Definition 7 (Local Minimal Fusion) We write $C \vdash^\sigma c$ iff:

$$\exists C' \subseteq C : (C'\sigma \vdash c\sigma \wedge \nexists \sigma' \subset \sigma : C'\sigma' \vdash c\sigma')$$

A local minimal fusion requires a subset C' of C (*locality* restriction) such that C' entails c whenever σ is applied. Also, we want σ to instantiate only those variables actually involved in the entailment of c – i.e. we require that no restriction σ' of σ suffices for the entailment.

To understand the motivations underlying the locality restriction, note that a substitution may be minimal w.r.t. a set of constraints C' , yet not minimal for a superset

$C \supset C'$, as the following example shows. Let $c = p(x)$, $C' = \{q(y), q(z) \vee s \rightarrow p(y)\}$ and $C = C' \cup \{s\}$. To obtain $C' \vdash c$, all the variables x, y, z must be instantiated to the same name. So, this fusion (call it σ_1) is minimal. Instead, to obtain $C \vdash c$, we can use another substitution σ_2 , which fuses x with y and neglects z , because the premise $q(z) \vee s$ can now be discharged through s . So, in this case σ_1 is *not* minimal, since $\sigma_2 \subset \sigma_1$. This phenomenon could, in principle, lead to unexpected behaviour. For instance, let:

$$\begin{aligned} P &= (x)(y)(z)(\text{fuse}_x c.R \mid R') \mid s \\ Q &= (x)(y)(z)(\text{fuse}_x c.R \mid R' \mid s) \end{aligned}$$

where R' is the parallel composition of the constraints C' . Let us drop for a while the locality restriction. Then, in P two minimal fusions would be applicable, depending on CLOSEFUSE being used at the top-level (σ_2) or not (σ_1). Instead, in Q only σ_2 would be possible. This clearly clashes with our intuition that P and Q should be equivalent, since Q is obtained from P through a scope extrusion. The locality restriction allows for recovering such equality. Indeed, σ_1 is minimal for $C' \subseteq C$, so it is also applicable for Q .

The consequence of locality is that inspecting any set of (locally) known contracts is enough to decide if a set of contracts leads to an agreement – it is not necessary to (globally) explore the whole system. To a local observer, fusing x, y, z may appear minimal. To a more informed observer, the same fusion may appear non-minimal (the minimal one fusing x, y). Both fusion are allowed by our semantics.

Example 4 (Handshaking) Recall from Sect. 1:

$$\begin{aligned} \text{Alice} &= (x) (\text{tell } b(x) \rightarrow a(x). \text{fuse}_x a(x). \text{lendAirplane}) \\ \text{Bob} &= (y) (\text{tell } a(y) \rightarrow b(y). \text{fuse}_y b(y). \text{lendBike}) \end{aligned}$$

A possible trace of $P = \text{Alice} \mid \text{Bob}$ is the following:

$$\begin{aligned} P &\xrightarrow{\tau} (x) (b(x) \rightarrow a(x) \mid \text{fuse}_x a(x). \text{lendAirplane}) \mid \\ &\quad (y) (a(y) \rightarrow b(y) \mid \text{fuse}_y b(y). \text{lendBike}) \\ &\xrightarrow{\tau} (n) (b(n) \rightarrow a(n) \mid \text{lendAirplane}\{n/x\} \mid \\ &\quad a(n) \rightarrow b(n) \mid \text{ask } b(n). \text{lendBike}\{n/y\}) \\ &\xrightarrow{\tau} (n) (b(n) \rightarrow a(n) \mid \text{lendAirplane}\{n/x\} \mid \\ &\quad a(n) \rightarrow b(n) \mid \text{lendBike}\{n/y\}) \end{aligned}$$

In the first two steps we fire the prefixes $\text{tell } b(x) \rightarrow a(x)$ and $\text{tell } a(y) \rightarrow b(y)$ through TELL , PARTAU , DEL . The third step is the crucial one. The prefix $\text{fuse}_x a(x)$ is fired through FUSE . Through CONSTR , PARFUSE , we discover the active constraint $c_a = b(x) \rightarrow a(x)$. We then use OPEN to obtain the action $(x)\{c_a\} \vdash_x^F a(x)$ for Alice. For Bob, we use CONSTR to discover $c_b = a(y) \rightarrow b(y)$, which we merge with the empty set of constraints obtained through IDLESUM ; we then use OPEN to get $(y)\{c_b\}$. At the top level, we then apply PARFUSE to deduce $(xy)\{c_a, c_b\} \vdash_x^F a(x)$. Finally, CLOSEFUSE fuses x and y to the fresh name n . Let $\sigma = \{n/x, n/y\}$ be

such fusion. It is easy to check that σ is a local minimal fusion, so $\{c_a, c_b\} \vdash^\sigma a(x)$. The instantiation transforms $\text{fuse}_y b(y)$ into $\text{ask } b(n)$, fired in the last step.

Example 5 (Unfair handshaking) To get further insights on the role played by contractual implication, consider an alternative handshaking, which makes *NO* use of \rightarrow .

$$\begin{aligned} \text{Alice}' &= (x) (\text{tell } a(x). \text{fuse}_x b(x). \text{lendAirplane}) \\ \text{Bob}' &= (y) (\text{tell } b(y). \text{fuse}_y a(y). \text{lendBike}) \end{aligned}$$

Since the handshaking is still performed, $P' = \text{Alice}' \mid \text{Bob}'$ behaves as $P = \text{Alice} \mid \text{Bob}$ in Ex. 4. Yet, P and P' behave quite differently in the presence of a third kid. Let:

$$\text{Carl} = (z) (\text{tell } c(z). \text{fuse}_z a(z). \text{lendCar})$$

The system $P' \mid \text{Carl}$ could lead to the execution of lendCar , since Carl can receive (through fuse) the promise $a(x)$ from Alice'. However, Alice' will be stuck waiting for a bike, so she will never respect her promise, and Carl will never obtain the expected airplane. In the system $P \mid \text{Carl}$, instead, Carl will not lend his car in vain. His fuse will be stuck, because the contract of Alice now requires a bike, and Carl does not provide it.

Example 6 (Principals) Consider an online market, where buyers and sellers trade items. The contract of a buyer n_B is to pay for an item, provided that some (still unknown) seller x_S promises to send it; dually, the contract of a seller n_S is to send an item, provided that some buyer y_B pays.

$$\begin{aligned} c_B &= n_B \text{ says } ((x_S \text{ says } \text{send}(x)) \rightarrow \text{pay}(x)) \\ c_S &= n_S \text{ says } ((y_B \text{ says } \text{pay}(y)) \rightarrow \text{send}(y)) \end{aligned}$$

A buyer first issues her contract c_B , then waits until discovering she has to pay, and eventually proceeds with the process B' . At this point, the buyer may either refuse to pay (process NoPay), or actually pay the item, by issuing a $\text{paid}(x)$. After the item has been paid, the buyer may wait for the item to be sent or open a dispute with the seller.

$$\begin{aligned} B &= (x)(x_S)(n_B) (\text{tell } c_B. \text{fuse}_x (n_B \text{ says } \text{pay}(x)). B') \\ B' &= \tau. \text{NoPay} + \tau. \text{tell } (n_B \text{ says } \text{paid}(x)). B'' \\ B'' &= \text{ask } (x_S \text{ says } \text{sent}(x)) + \tau. \text{tell } (n_B \text{ says } \text{dispute}(x)) \end{aligned}$$

The behaviour of the seller n_S is dual.

$$\begin{aligned} S &= (y)(y_B)(n_S) (\text{tell } c_S. \text{fuse}_y (n_S \text{ says } \text{send}(y)). S') \\ S' &= \tau. \text{NoSend} + \tau. \text{tell } (n_S \text{ says } \text{sent}(y)). S'' \\ S'' &= \text{ask } (y_B \text{ says } \text{pay}(y)) + \tau. \text{tell } (n_S \text{ says } \text{dispute}(y)) \end{aligned}$$

An handshaking is reached through CLOSEFUSE . The fusion is $\sigma = \{m/x, m/y, n_S/x_S, n_B/y_B\}$, where m is a fresh name.

To automatically resolve disputes, a judge J can enter a session initiated between a buyer and a seller, provided that

a dispute has been opened, and either the obligations pay or send have been inferred. This is done through the $\text{ask}_{\vec{x}}$ primitive, where $\vec{x} = \{z, x_S, y_B\}$. This binds the variable z to the session identifier m , x_S to the actual name of the seller (n_S), and y_B to the actual name of the buyer (n_B).

$$J = (\vec{x})(\text{ask}_{\vec{x}}(y_B \text{ says pay}(z) \wedge x_S \text{ says dispute}(z)). \\ \text{check } \neg(y_B \text{ says paid}(z)). \text{jail}(y_B) \\ | \text{ask}_{\vec{x}}(x_S \text{ says send}(z) \wedge y_B \text{ says dispute}(z)). \\ \text{check } \neg(x_S \text{ says sent}(z)). \text{jail}(x_S))$$

If the obligation $\text{pay}(z)$ is found, but the item has not been actually paid then the buyer is convicted (modelled by $\text{jail}(y_B)$, not further detailed). Similarly, if the obligation $\text{send}(z)$ has not been supported by a corresponding $\text{sent}(z)$, then the seller is convicted.

Encodings. We have studied the expressive power of our contract calculus by encoding several concurrency idioms, among which semaphores, Linda [23], the π -calculus and graph rewriting. See [5] for all the details.

4 Related Work

The complexity of real-world scenarios, where several concepts like principals, contracts, authorizations, duties, delegation, mandates, regulations, *etc.* are inextricably intermingled, have led to a steady flourishing of new logics over the years. These take inspiration and extend e.g. classical [17], modal [16], deontic [34, 22], default [24] and defeasible logics [25]. We think none of these logics, including our PCL, captures *all* the facets of contracts. Each of these logics is designed to represent some particular aspect of contracts, e.g. obligations, permissions and prohibitions in deontic logics, violation of contracts in default and defeasible logics, and agreement in our contract logic. We argue that, since these aspects are orthogonal, it is possible to extend PCL with features from some of these logics.

The motivations underlying our logic seem related to those introduced in [3] to compose assume-guarantee specifications [2]. The idea is that a system will give some guarantee M_1 about its behaviour, provided that the environment it operates within will behave according to some assumption M_2 , and *vice versa*. This is rendered in [3] as the judgment $(M_1 \rightarrow M_2) \wedge (M_2 \rightarrow M_1) \vdash M_1 \wedge M_2$. However, since \rightarrow is the usual intuitionistic implication, this judgment (not valid in IPC) only holds in particular models, where e.g. M_1, M_2 must be interpreted as safety properties. Our approach is different: we make the above judgement valid by using \Rightarrow instead of \rightarrow . We then develop a decidable proof theory (exploited to design our calculus of contracting processes) while abstracting from the specific models. Actually, finding sound and complete models for contracts

seems to be hard, e.g. [31] shows the impossibility of devising, in some lattice-based models, a sound and complete set of compositional rules for circular assume-guarantee.

Our research seems also related to foundational research on authorization logics for distributed systems [1, 21, 30]. A crucial difference is that, while authorizations logics are focussed on deciding, given a bunch of logical authorization assertions, *if* a principal is allowed to perform some action, in our contract logic we are also concerned with discovering *what* that principal has to promise in return.

In our model of contracts we have abstracted from most of the implementation issues. For instance, in insecure environments populated by attackers, the operation of exchanging contracts requires particular care. This is related to the problem of establishing *common knowledge* in distributed systems [26]. A trusted third party might then be in order to make the contract exchange fair. We expect to apply standard techniques for guaranteeing non-repudiation [28, 39], fair exchange [4], and contract signing [15, 27]. Note that our logic works at a higher abstraction level than the protocol level, by focussing on how the exchanged messages lead to an handshaking between the parties.

Contracts are modelled as processes in [12, 13], to specify the interaction behaviour of clients and services. A client contract complies with a service contract if any possible interaction will always succeed. The crux is how to define (and decide) a subcontract relation, that allows for safely substituting services without affecting the compliance with their clients. Our contracts could be seen as a declarative underspecified description of which behavioural contracts are an implementation. Behavioural contracts seem more rigid than ours, as they precisely fix the order in which the actions must be performed. Even though in some cases this may be desirable, many real-world contracts allow for a more liberal way of constraining the involved parties (e.g., “I will pay before the deadline”). While the crucial notion in [12] is *compatibility* (which results in a yes/no output), we focus on the inferring the *obligations* that arise from a set of contracts. This provides a fine-grained quantification of the reached agreement, e.g. we may identify who is responsible of a contract violation.

Negotiation and service-level agreement are dealt with in cc-pi [10], a calculus combining features from concurrent constraints and name passing. As in the π -calculus, synchronization is channel-based: it only happens between two processes sharing a name. Synchronization fuses two names, similarly to the fusion calculus and ours. A main difference between cc-pi and our calculus is that in cc-pi only two parties may simultaneously reach an agreement, while our fuse allows for simultaneous multiparty agreements. Also, in our calculus the parties involved in an agreement do not have to share a pre-agreed name. This is useful for modelling scenarios where a contract can be accepted

by any party meeting the required terms (see e.g. Ex. 6).

Our contracts could be exploited to enhance the compensation mechanism of long-running transactions [7, 8, 11]. There, a long transaction is partitioned into a sequence of smaller ones, each one associated with a *compensation*, to be run upon failures of the standard execution [20]. While in long-running transactions clients have little or no control on the compensations (they are specified by the designer), in our approach clients can use contracts to select those services offering the desired compensation. In [9], cc-pi is extended with rules for handling transactions. This gives to the client more control on the choice of compensations. The differences w.r.t. our calculus noted above still apply.

The *reputation* of the principal issuing a contract could be influential in deciding whether accepting such contract or not. Reputation systems which evaluate the past behaviour of principals, e.g. [29] can then be used to make our contracts more expressive. Time is another feature that may arise while modelling contracts. Temporal extensions of our logics, e.g. like [18], will allow to check whether a promise is violated in a given trace (e.g. the deadline passed).

References

- [1] M. Abadi, M. Burrows, B. Lampson, and G. Plotkin. A calculus for access control in distributed systems. *ACM TOPLAS*, 4(15):706–734, 1993.
- [2] M. Abadi and L. Lamport. Composing specifications. *ACM TOPLAS*, 15(1):73–132, 1993.
- [3] M. Abadi and G. D. Plotkin. A logical view of composition. *Theoretical Computer Science*, 114(1):3–30, 1993.
- [4] N. Asokan, V. Shoup, and M. Waidner. Optimistic fair exchange of digital signatures. In *EUROCRYPT*, 1998.
- [5] M. Bartoletti and R. Zunino. A calculus of contracting processes. Technical Report DISI-09-056, Univ. Trento, 2009.
- [6] M. Bartoletti and R. Zunino. A logic for contracts. Technical Report DISI-09-034, Univ. Trento, 2009. (Nov 2009 revision).
- [7] L. Bocchi, C. Laneve, and G. Zavattaro. A calculus for long running transactions. In *Proc. FMOODS*, 2003.
- [8] R. Bruni, H. C. Melgratti, and U. Montanari. Theoretical foundations for compensations in flow composition languages. In *Proc. POPL*, 2005.
- [9] M. G. Buscemi and H. C. Melgratti. Transactional service level agreement. In *Proc. TGC*, 2007.
- [10] M. G. Buscemi and U. Montanari. CC-Pi: A constraint-based language for specifying service level agreements. In *Proc. ESOP*, 2007.
- [11] M. J. Butler, C. A. R. Hoare, and C. Ferreira. A trace semantics for long-running transactions. In *25 Years Communicating Sequential Processes*, 2004.
- [12] G. Castagna, N. Gesbert, and L. Padovani. A theory of contracts for web services. *ACM TOPLAS*, 31(5), 2009.
- [13] G. Castagna and L. Padovani. Contracts for mobile processes. In *Proc. CONCUR*, 2009.
- [14] M. Coppo and M. Dezani-Ciancaglini. Structured communications with concurrent constraints. In *Proc. TGC*, 2008.
- [15] V. Cortier, R. Küsters, and B. Warinschi. A cryptographic model for branching time security properties – the case of contract signing protocols. In *Proc. ESORICS*, 2007.
- [16] A. Daskalopulu and T. Maibaum. Towards electronic contract performance. In *Proc. DEXA*, 2001.
- [17] H. Davulcu, M. Kifer, and I. Ramakrishnan. CTR-S: A logic for specifying contracts in semantic web services. In *Proc. WWW*, 2004.
- [18] H. DeYoung, D. Garg, and F. Pfenning. An authorization logic with explicit time. In *Proc. CSF*, 2008.
- [19] M. Fairtlough and M. Mendler. Propositional lax logic. *Information and Computation*, 137(1):1–33, 1997.
- [20] H. Garcia-Molina and K. Salem. Sagas. In *SIGMOD Conference*, 1987.
- [21] D. Garg and M. Abadi. A modal deconstruction of access control logics. In *Proc. FoSSaCS*, 2008.
- [22] J. Gelati, A. Rotolo, G. Sartor, and G. Governatori. Normative autonomy and normative co-ordination: Declarative power, representation, and mandate. *Artificial Intelligence and Law*, 12(1-2):53–81, 2004.
- [23] D. Gelernter. Generative communication in Linda. *ACM TOPLAS*, 7(1), 1985.
- [24] G. K. Giannakis and A. Daskalopulu. The representation of e-contracts as default theories. In *New Trends in Applied Artificial Intelligence*, 2007.
- [25] G. Governatori. Representing business contracts in RuleML. *Int. J. of Cooperative Information Systems*, 14(2-3), 2005.
- [26] J. Y. Halpern and Y. Moses. Knowledge and common knowledge in a distributed environment. *J. ACM*, 37(3), 1990.
- [27] D. Kähler and R. Küsters. Constraint solving for contract-signing protocols. In *Proc. CONCUR*, 2005.
- [28] S. Kremer, O. Markowitch, and J. Zhou. An intensive survey of fair non-repudiation protocols. *Computer Communications*, 25, 2002.
- [29] K. Krukow, M. Nielsen, and V. Sassone. A logical framework for history-based access control and reputation systems. *Journal of Computer Security*, 16(1):63–101, 2008.
- [30] N. Li, B. N. Grosz, and J. Feigenbaum. Delegation logic: A logic-based approach to distributed authorization. *ACM TISSEC*, 6(1):128–171, 2003.
- [31] P. Maier. Compositional circular assume-guarantee rules cannot be sound and complete. In *Proc. FoSSaCS*, 2003.
- [32] J. Parrow and B. Victor. The fusion calculus: Expressiveness and symmetry in mobile processes. In *Proc. LICS*, 1998.
- [33] F. Pfenning. Structural cut elimination - intuitionistic and classical logic. *Information and Computation*, 157, 2000.
- [34] C. Prisacariu and G. Schneider. A formal language for electronic contracts. In *Proc. FMOODS*, 2007.
- [35] V. Saraswat, P. Panangaden, and M. Rinard. Semantic foundations of concurrent constraint programming. In *Proc. POPL*, 1991.
- [36] R. Statman. Intuitionistic propositional logic is polynomial-space complete. *Theoretical Computer Science*, 9, 1979.
- [37] The PCL web site. <http://www.disi.unitn.it/~zunino/PCL>.
- [38] M. Wirsing et al. SENSORIA process calculi for service-oriented computing. In *Proc. TGC*, 2006.
- [39] J. Zhou. *Non-repudiation in Electronic Commerce*. Artech House, 2001.

A Selected Cases for Cut elimination

Case ZERO / PREPOST

$$\frac{\frac{\frac{D0}{\Gamma \vdash q} \text{ZERO}}{\Gamma \vdash p \rightarrow q} \quad \frac{\frac{\frac{D1}{\Gamma, p \rightarrow q, a \vdash p} \quad \frac{D2}{\Gamma, p \rightarrow q, q \vdash b} \text{PREPOST}}{\Gamma, p \rightarrow q \vdash a \rightarrow b}}{\Gamma \vdash a \rightarrow b} \text{CUT}}{\Gamma \vdash a \rightarrow b} \text{CUT} \Rightarrow \frac{\frac{\frac{D0}{\Gamma \vdash q} \quad \frac{\frac{\overline{\Gamma, q \vdash q} \text{ID}}{\Gamma, q \vdash p \rightarrow q} \text{ZERO}}{\Gamma, q \vdash b} \text{CUT}}{\Gamma \vdash b} \text{CUT}}{\Gamma \vdash a \rightarrow b} \text{ZERO}$$

Case ZERO / FIX

$$\frac{\frac{\frac{D0}{\Gamma \vdash q} \text{ZERO}}{\Gamma \vdash p \rightarrow q} \quad \frac{\frac{\frac{D1}{\Gamma, p \rightarrow q, a \vdash p} \quad \frac{D2}{\Gamma, p \rightarrow q, q \vdash a} \text{FIX}}{\Gamma, p \rightarrow q \vdash a} \text{CUT}}{\Gamma \vdash a} \text{CUT}}{\Gamma \vdash a} \text{CUT} \Rightarrow \frac{\frac{\frac{D0}{\Gamma \vdash q} \quad \frac{\frac{\overline{\Gamma, q \vdash q} \text{ID}}{\Gamma, q \vdash p \rightarrow q} \text{ZERO}}{\Gamma, q \vdash a} \text{CUT}}{\Gamma \vdash a} \text{CUT}}{\Gamma \vdash a} \text{CUT}$$

Case PREPOST / FIX, assuming $(p \rightarrow q) \in \Gamma$

$$\frac{\frac{\frac{D0}{\Gamma, a \vdash p} \quad \frac{D1}{\Gamma, q \vdash b} \text{PREPOST}}{\Gamma \vdash a \rightarrow b} \quad \frac{\frac{\frac{D2}{\Gamma, a \rightarrow b, r \vdash a} \quad \frac{D3}{\Gamma, a \rightarrow b, b \vdash r} \text{FIX}}{\Gamma, a \rightarrow b \vdash r} \text{CUT}}{\Gamma(p \rightarrow q) \vdash r} \text{CUT} \Rightarrow \frac{\frac{\hat{D}_0}{\Gamma, r \vdash p} \quad \frac{\hat{D}_1}{\Gamma, q \vdash r} \text{FIX}}{\Gamma \vdash r} \text{FIX}$$

$$\text{where } \hat{D}_0 = \frac{\frac{\frac{D0+}{\Gamma, r, a \vdash p} \quad \frac{D1+}{\Gamma, r, q \vdash b} \text{PREPOST}}{\Gamma, r \vdash a \rightarrow b} \quad \frac{D2}{\Gamma, r, a \rightarrow b \vdash a} \text{CUT}}{\Gamma, r \vdash a} \text{CUT} \quad \frac{D0+}{\Gamma, r, a \vdash p} \text{CUT}$$

$$\text{and } \hat{D}_1 = \frac{\frac{\frac{D1}{\Gamma, q \vdash b} \quad \frac{\frac{\frac{D0+}{\Gamma, q, b, a \vdash p} \quad \frac{D1+}{\Gamma, q, b, q \vdash b} \text{PREPOST}}{\Gamma, q, b \vdash a \rightarrow b} \quad \frac{D3+}{\Gamma, q, b, a \rightarrow b \vdash r} \text{CUT}}{\Gamma, q, b \vdash r} \text{CUT}}{\Gamma, q \vdash r} \text{CUT}$$

Case FIX / any rule, assuming $(p \rightarrow q) \in \Gamma$

$$\frac{\frac{\frac{D0}{\Gamma, a \vdash p} \quad \frac{D1}{\Gamma, q \vdash a} \text{FIX}}{\Gamma \vdash a} \quad \frac{D2}{\Gamma, a \vdash b} \text{CUT}}{\Gamma(p \rightarrow q) \vdash b} \text{CUT} \Rightarrow \frac{\frac{\frac{\frac{D1+}{\Gamma, b, q \vdash a} \quad \frac{D0+}{\Gamma, b, q, a \vdash p} \text{CUT}}{\Gamma, b, q \vdash p} \text{FIX}}{\Gamma, b \vdash p} \text{ID}}{\Gamma, b \vdash p} \text{FIX}}{\Gamma \vdash b} \text{FIX} \quad \frac{\frac{D1}{\Gamma, q \vdash a} \quad \frac{D2+}{\Gamma, q, a \vdash b} \text{CUT}}{\Gamma, q \vdash b} \text{CUT}$$

Our cut metric is the following. We associate to each cut the weight $(p, h(D_0), h(D_1))$ where p is the cut formula, and $h(D_i)$ denotes the height of the subderivation D_i . The structural ordering of formulae and the usual one on naturals induce the lexicographical ordering of weights. This is a well-founded ordering, and each reduction above replaces a cut with *lighter* ones, only.