

Directed Deadline Obligations in Agent-Based Business Contracts

Henrique Lopes Cardoso and Eugénio Oliveira

LIACC, DEI / Faculdade de Engenharia, Universidade do Porto
R. Dr. Roberto Frias, 4200-465 Porto, Portugal
{hlc,eco}@fe.up.pt

Abstract. There are B2B relationships that presume cooperation in contract enactment. This issue should be taken into account when modeling, for computational handling, contractual commitments through obligations. Deadline obligations have been modeled by considering that reaching the deadline without compliance brings up a violation. When modeling commitments in business contracts, directed obligations have been studied for identifying two agents: the obligation's bearer and the counterparty, who may claim for legal action in case of non-compliance. We argue in favor of a directed deadline obligation approach, taking inspiration on international legislation over trade procedures. Our proposal to model contractual obligations is based on authorizations granted in specific states of an obligation lifecycle model, which we formalize using temporal logic and implement in a rule-based system. The performance of a contractual relationship is supported by a model of flexible deadlines, which allow for further cooperation between autonomous agents. As a result, the decision-making space of agents concerning contractual obligations is enlarged and becomes richer. We discuss the issues that agents should take into account in this extended setting.

1 Introduction

In cooperative B2B Virtual Organizations, agents (representing different enterprises) share their own competences and skills in a regulated way, through commitments expressed as norms in contracts. The importance of successfully proceeding with business demands for flexibility of operations: agents should try to facilitate the compliance of their partners. This common goal of conducting a multiparty business is based on the fact that group success also benefits each agent's private goals. These goals are not limited to the ongoing business relationship, but also concern future opportunities that may arise.

While addressing this problem with norms and multi-agent systems, we find that many approaches to normative multi-agent systems are abstracted away from their potential application domain. As such, deontic operators are often

taken to have a universal semantics. For instance, deadline obligations are violated if the obliged action or state is not obtained until the deadline is reached.

We argue that in some domains – such as in business contracts – such an approach is not desirable. For instance, the United Nations Convention on Contracts for the International Sale of Goods (CISG) [1] establishes what parties may do in case of deadline violations. In some cases they are allowed to fulfill their obligations after the deadline (Article 48), or even to extend the deadlines with the allowance of their counterparties. Furthermore, a party may extend his counterparty’s deadlines (Articles 47 and 63), which denotes a flexible and even cooperative facet of trade contracts.

In this paper we propose a different approach (in comparison with [2][3][4][5]) to the use of obligations in MAS in the domain of business contracts. Following a cooperative business performance posture, we argue that obligations should be directed, and that deadlines should be flexible. We start by reviewing, in section 2, the most typical variations regarding the formalization of obligations, after which we propose an approach based on directed obligations with deadlines. The flexibility required when handling temporal restrictions of obligations is addressed in section 3. The proposed approach is based on authorizations, and we present a lifecycle for directed obligations with temporal restrictions. In section 4 we investigate the decision-making process of agents concerning authorizations. Implementation of the proposed model in a rule-based system is discussed in section 5. Section 6 discusses related work and section 7 concludes.

2 Contractual Obligations

The use of norms in MAS makes use of the well-known deontic operators of obligation, permission and prohibition [6]. In theoretical deontic logic approaches, these operators are sometimes used to represent abstract general principles (e.g. it is forbidden to kill). In more applied research, deontic operators are ascribed either to roles or to particular agents in a system; e.g. $O_b(f)$ indicates that agent b is obliged to bring about fact f (a state of affairs or an action) – in this case agent b is said to be the *bearer* of the obligation.

Also, deontic operators are often made conditional and time constrained. Considering obligations, the conditional aspect has taken two different perspectives: conditional obligations of the form $O_b(f/s)$, meaning that agent b is obliged to bring about f when situation s arises; and conditional norms of the form $s \rightarrow O_b(f)$, meaning that if s then b is obliged to bring about f . As for the temporal aspect of deontic operators, deadlines (either time references or more generally defined as states of affairs) are typically employed for stipulating the validity of the operator: $O_b(f, d)$ is a *deadline obligation* indicating that agent b is obliged to bring about f before d .

We will base the following discussion on the *obligation* deontic operator, as it is the most important operator to represent trade relationships in B2B contracts. The meaning of deontic operators has been studied, mainly regarding the use of

deadlines (e.g. [2]). Regarding deadline obligations, the usual approach to their semantics is to consider the following entailments¹:

- $O_b(f, d) \wedge (f B d) \models \text{Fulf}_b(f, d)$ — If the fact to bring about occurs before the deadline, the agent has fulfilled his obligation.
- $O_b(f, d) \wedge (d B f) \models \text{Viol}_b(f, d)$ — If the deadline occurs before the fact to bring about, the agent has violated his obligation.

The introduction of *Fulf* and *Viol* enables reasoning about the respective situations. The implementation of this semantics using forward-chaining rules has been studied in [3]. Although intuitive, this semantics is quite rigid in that violations are all defined in a universal way (discounting the fact that different norms can respond to violations in different ways).

The analysis of contracts brings into discussion the notion of *directed obligations* [8]. Obligations are seen as directed from a *bearer* (responsible for fulfilling the obligation) to a *counterparty*. Some authors [4] define the very notion of *contractual obligation* as an obligation with an “obligor” (bearer) and an “obligee” (counterparty). The relationship between these two roles in a directed obligation has been studied, giving rise to two different theories. The *benefit theory* promotes the fact that the counterparty of an obligation is intended to benefit from its fulfillment (see [8] for a benefit theory perspective of directed obligations). A more relevant approach in which contracts are concerned – the *claimant theory* – takes the stance that obligations are interpreted as claims from counterparties to bearers (see [5] for a claimant theory support).

In general, claimant approaches are based on the following definition for directed obligation (adapted from [5]): $O_{b,c}(f) =_{def} O_b(f) \wedge (\neg f \Rightarrow P_c(la_b))$. A directed obligation from agent b towards agent c to bring about f means that b is obliged to bring about f and if b does not bring about f then c is *permitted* to initiate legal action against b . The concept of legal action is rather vague. A similar approach is taken in [9], where agent c is said to be *authorized* to repair the situation in case b does not fulfill his obligation. Repair actions include demanding further actions from b ; e.g., c may demand compensation for damages. It is interesting to note that such definitions are careful enough to base the claims of the counterparty on the *non-fulfillment* of the obligation, not on its violation. In fact, these definitions do not include deadlines, which are the basis for violation detection. Another significant issue is the discretionary nature of the counterparty’s reaction (he is permitted or authorized), instead of an automatic response based on the non-fulfillment of the bearer².

¹ In the following formulae we will follow linear temporal logic (LTL) [7], with a discrete time model. Let $x = (s_0, s_1, s_2, \dots)$ be a timeline, defined as a sequence of states s_i . The syntax $x \models p$ reads that p is true in timeline x . We write x^k to denote state s_k of x , and $x^k \models p$ to mean that p is true at state x^k . We use a weak version of the *before* LTL operator B , where q is not mandatory: $x \models (p B q)$ iff $\exists_j (x^j \models p \wedge \forall_{k < j} (x^k \models \neg q))$.

² As in automatic violation detection approaches based on deadlines, complemented with the definition of violation reaction norms.

We propose the use of *directed deadline obligations* as the basis for defining contractual obligations: $O_{b,c}(f, d)$. In section 3 we describe a model for flexible obligation violation, based on the principle that the deadline is meant to indicate when the counterparty is authorized to react to the non-fulfillment of an obligation directed to him. A possible reaction is to declare the obligation as violated, but there are other means to settle the matter, to the benefit of both involved parties. An extension of directed (contractual) obligations with temporal restrictions is also introduced in [4], but that approach is based on a rigid model of violations, in that they are automatically obtained at the deadline.

2.1 Directed Deadline Obligations

Our proposal combines directed [5][8] and deadline [2] obligations. Although this has been done in the past (e.g. [4]), in our approach deadlines have a distinct role in the definition of obligations. In section 3 we detail such a role.

Directed deadline obligations take the form $O_{b,c}(f, d)$, meaning that agent b is obliged towards agent c to bring about f before d . We do not make obligations conditional (as in [4]), because we assume they are obtained from conditional norms: rules prescribing obligations when certain situations arise.

We consider that if fact f is not yet the case when deadline d arises, the obligation is not yet violated, but is in a state where the counterparty is authorized to take some action. We emphasize the case for a *deadline violation* (as opposed to obligation violation). This comprises a flexible approach to handling non-ideal situations: each deadline violation is different, as each may have a different impact on the ongoing business, and each occurs between a specific pair of agents with a unique trust relationship.

Some evidence from the CISG convention [1] led us to this approach:

Article 48: (1) [...] the seller may, even after the date for delivery, remedy at his own expense any failure to perform his obligations, if he can do so without unreasonable delay [...]; (2) If the seller requests the buyer to make known whether he will accept performance and the buyer does not comply with the request within a reasonable time, the seller may perform within the time indicated in his request. [...]

This means that even though a deadline has been violated, the bearer may still be entitled to fulfill *the same* obligation. This kind of delay is also called a *grace period*: a period beyond a due date during which an obligation may be met without penalty or cancellation.

Figure 1 illustrates the intuitive semantics of a directed deadline obligation. The shaded area represents the period of time within which the achievement of f will certainly bring a fulfillment of the obligation. The region to the right of d indicates that counterparty c is entitled to react if f is not accomplished; however, as long as no reaction is taken, b can still fulfill his obligation.

Therefore, a deadline violation brings a counterparty authorization. Authorizations are taken into account in the normative system by having rules and




Fig. 1. Directed obligation with deadline

norms that are based on the materialization of such authorizations. The available options are discussed in section 3.

2.2 Livelines and Deadlines

The deadline approach is often taken to be appropriate for specifying temporal restrictions on obligations. However, in certain cases a time window should be provided. In international trade transactions, for instance, storage costs may be relevant. Also, perishable goods should be delivered only when they are needed, not before. This is why in CISG [1] we have:

Article 52: (1) If the seller delivers the goods before the date fixed, the buyer may take delivery or refuse to take delivery.

Therefore, anticipated fulfillments are not always welcome. We find it necessary to include a variation of directed deadline obligations, to which we add a *liveline*: a time reference after which the obligation should be fulfilled. In this case we have $O_{b,c}(f, l, d)$: agent b is obliged towards agent c to bring about f between l (a liveline) and d (a deadline). Figure 2 illustrates the intuitive semantics of this kind of obligation. The shaded area represents the period of time within which the achievement of f will certainly bring a fulfillment of the obligation. If f is accomplished before l , however, it may be the case that c is not willing to accept such a fulfillment, or at least that he may not be happy about it – the region to the left of l entitles c to react if f is accomplished. The region to the right of d is as with (simple) directed deadline obligations.




Fig. 2. Directed obligation with liveline and deadline

We escape from an approach with a fixed time reference for obligation fulfillment (an obligation for bringing about f at time t), which would be suggested

by the term “date fixed” in CISG’s Article 52 transcription above. We find it more convenient to define a fixed date as an interval, say, from the beginning till the end of a specific date³.

3 Managing Liveline and Deadline Violations

After we have advocated, in the preceding section, a counterparty authorization approach to deadline violations, in this section we address the issue of what kind of actions the counterparty may take in such situations, and what are their effects on the obligation whose deadline has been violated. The same accounts to directed obligations with both livelines and deadlines.

The successful enactment of a contract is dependent on the need to make contractual provisions performable in a flexible way. The importance of having flexible trade procedures is apparent, once again, in the CISG convention [1]:

Article 47: (1) The buyer may fix an additional period of time of reasonable length for performance by the seller of his obligations.

Article 63: (1) The seller may fix an additional period of time of reasonable length for performance by the buyer of his obligations.

These articles emphasize, once more, the need for flexible deadlines. Note that the counterparty’s benevolence on conceding an extended deadline to the bearer does not prescribe a new obligation; instead, *the same* obligation may be fulfilled within a larger time window. Furthermore, it is also in the counterparty’s best interest that this option is available, given the importance of reaching success in the performance of the contract.

In some other cases, a party may decide that the non-fulfillment of an obligation should be handled in a more strict way. The CISG convention specifies conditions for cancelling a contract in case of breach:

Article 49: (1) The buyer may declare the contract avoided: (a) if the failure by the seller to perform any of his obligations [...] amounts to a fundamental breach of contract; [...]; (2) However, in cases where the seller has delivered the goods, the buyer loses the right to declare the contract avoided unless he does so: (a) in respect of late delivery, within a reasonable time after he has become aware that delivery has been made; [...]

Article 64: (1) The seller may declare the contract avoided: (a) if the failure by the buyer to perform any of his obligations [...] amounts to a fundamental breach of contract; [...]; (2) However, in cases where the buyer has paid the price, the seller loses the right to declare the contract avoided unless he does so: (a) in respect of late performance by the buyer, before the seller has become aware that performance has been rendered; [...]

³ This is actually a matter of time granularity.

These articles allow contract termination in both non-performance and late performance cases. However, the second case is limited to the awareness of the offended party.

From these excerpts we can distinguish two types of reactions to non-fulfillments: a smoother one (from articles 47, 48 and 63), in which parties are willing to recover from an initial failure to conform to an obligation; and a stricter one (articles 49 and 64), where the failure is not self-containable anymore. Based on these options, we propose a model for a directed deadline obligation lifecycle.

3.1 Authorizations on Violations

Following the discussion above, we identify the possible states for an obligation, together with the elements we shall use to signal some of those states (when obtained, these elements are supposed to persist over time):

- *inactive*: the obligation is not yet in effect, but will eventually be prescribed by a norm;
- *active*: the obligation was prescribed by a norm – $O_{b,c}(f, d)$ or $O_{b,c}(f, l, d)$
- *pending*: the obligation may be fulfilled from now on;
- *liveline violation*: the fact being obliged has been brought ahead of time – $LViol_{b,c}(f, l, d)$
- *deadline violation*: the fact being obliged should have been brought already – $DViol_{b,c}(f, d)$ or $DViol_{b,c}(f, l, d)$
- *fulfilled*: the obligation was fulfilled – $Fulf_{b,c}(f, d)$ or $Fulf_{b,c}(f, l, d)$
- *violated*: the obligation was violated and cannot be fulfilled anymore – $Viol_{b,c}(f, d)$ or $Viol_{b,c}(f, l, d)$

Starting with the simpler case of directed deadline obligations, we identify the (absolute) fulfillment case:

- $O_{b,c}(f, d) \wedge (f B d) \models Fulf_{b,c}(f, d)$

Then we state the consequence of reaching a deadline with no achievement of the obligated fact:

- $O_{b,c}(f, d) \wedge (d B f) \models DViol_{b,c}(f, d)$

Note that, differently from the usual approach, we set the obligation to have a violated deadline – $DViol_{b,c}(f, d)$ – but not to be violated in itself.

The counterparty's reaction to a deadline violation will only change the obligation's state if the option is to deem the obligation as violated, by *denouncing* this situation. For this we introduce the element $Den_{c,b}(f, d)$, which is a denounce from agent c towards agent b regarding the failure of the latter to comply with his obligation to bring about f before d . Since we consider the achievement of facts to be common knowledge, a party may only denounce the non-fulfillment of an obligation while that obligation is not fulfilled yet⁴:

⁴ This is a simplification of what articles 49 and 64 of CISG suggest.

- $DViol_{b,c}(f, d) \wedge (f B Den_{c,b}(f, d)) \models Fulf_{b,c}(f, d)$
- $DViol_{b,c}(f, d) \wedge (Den_{c,b}(f, d) B f) \models Viol_{b,c}(f, d)$

Figure 3 illustrates, by means of a state transition diagram, the lifecycle of directed deadline obligations. We take obligations as being prescribed from conditional norms; the confirmation of the norm’s condition will change the prescribed obligation’s state from inactive to active. The obligation is also automatically pending, since it may be legitimately fulfilled right away.




Fig. 3. Lifecycle of a directed deadline obligation




Fig. 4. Lifecycle of a directed obligation with liveline and deadline

Figure 4 contains the state transition diagram for directed obligations with livelines and deadlines. In this case, the obligation will only be pending when l arises, since only then it may be fulfilled in a way that is compliant with the terms of the contract. We define the following relations:

- $O_{b,c}(f, l, d) \wedge (f B l) \models LViol_{b,c}(f, l, d)$
- $LViol_{b,c}(f, l, d) \wedge (l B Den_{c,b}(f, l, d)) \models Fulf_{b,c}(f, l, d)$
- $LViol_{b,c}(f, l, d) \wedge (Den_{c,b}(f, l, d) B l) \models Viol_{b,c}(f, l, d)$
- $O_{b,c}(f, l, d) \wedge (l B f) \wedge (f B d) \models Fulf_{b,c}(f, l, d)$
- $O_{b,c}(f, l, d) \wedge (d B f) \models DViol_{b,c}(f, l, d)$
- $DViol_{b,c}(f, l, d) \wedge (f B Den_{c,b}(f, l, d)) \models Fulf_{b,c}(f, l, d)$
- $DViol_{b,c}(f, l, d) \wedge (Den_{c,b}(f, l, d) B f) \models Viol_{b,c}(f, l, d)$

We have now two kinds of temporal violations: liveline violations of the form $LViol_{b,c}(f, l, d)$ and deadline violations of the form $DViol_{b,c}(f, l, d)$. In both cases a denounce may establish the obligation as violated, if issued before l or f , respectively.

3.2 Smoother Authorizations on Violations

The diagrams in figures 3 and 4 only include events that produce a change in an obligation's state. The denouncement of the non-fulfillment of an obligation, making it violated and consequently not fulfillable any longer, denotes a situation in which a bearer's attempt to fulfill the obligation will no longer be significant to the counterparty, and thus a consummated violation should be handled according to applicable norms. These may bring sanctions, further obligations or ultimately a contract cancellation, as in articles 49 and 64 of CISG.

In order to accommodate less strict situations (see articles 47, 48 and 63 of CISG), we consider that in liveline and deadline violation states, while the obligation can still be fulfilled, the counterparty may react to the non-ideal situation. These possibilities are not illustrated in figures 3 and 4, since they do not bring state changes. For instance, in international trade transactions storage costs may be relevant. The counterparty may therefore be authorized to demand for payment of storage costs from an early compliant bearer. Another example for the deadline violation case:

Article 78: If a party fails to pay the price or any other sum that is in arrears, the other party is entitled to interest on it [...]

While obligation state transitions are processed with appropriate rules (including rules that take denounces into account), authorizations expressing the counterparty's right to demand for compensation are handled by the system through appropriate norms, which may be defined in a contract basis.

4 Decision-Making on Directed Deadline Obligations

The authorization approach described above enriches the decision-making space of agents concerning norms. Since commitments can be violated, agents (as human delegates) may decide whether to fulfill them or not. Furthermore, because the violation state is determined by the counterparty's choice to denounce this situation, both parties associated with a directed obligation are in a position to decide over it after the deadline.

In order to model the decision making process, we need to assess each agent's valuations on the obligation states and facts they are able to bring about. We will write $v_a(f)$ and $v_a(S)$ to denote the valuation agent a makes of fact f or state S , respectively (similarly to the valuation model used in [10]). When valuating an obligation's state (namely a fulfillment or a violation), agents should take into account two different sorts of effects. First, since an obligation is taken to be a part of a wider contract that should benefit all participants, the obligation cannot be taken in isolation, as its fulfillment or violation may trigger further commitments. Second, an agent's reputation is affected by whether or not he stands for his commitments. In the following we assume that an agent is capable of anticipating and evaluating the consequences of his actions within a contract.

For an obligation $O_{b,c}(f, d)$ we have the following valuation constraints for b :

$v_b(O_{b,c}(f, d)) < 0$: an obligation is a burden to its bearer
 $v_b(f) < v_b(O_{b,c}(f, d))$: there is a heavier cost associated with bringing about f
 $v_b(Fulf_{b,c}(f, d)) > 0$: b gains from fulfilling his obligation
 $v_b(Viol_{b,c}(f, d)) < 0$: b loses from violating his obligation

The notions of gain and loss for the bearer extend to outside this obligation. For instance, fulfilling an obligation may bring an entitlement (a new obligation where the bearer becomes the counterparty). Violating an obligation will potentially bring penalties to the bearer, hence the negative valuation. In both cases, the reputation of agent b is affected (positively or negatively). Unlike in [10], we do not impose that $v_b(Viol_{b,c}(f, d)) < v_b(f) + v_b(Fulf_{b,c}(f, d))$. An agent may be able to exploit a contract flaw by considering that in a specific situation he is better off violating his obligation than fulfilling it. Of course that even if the above condition holds, agent b may still choose to violate his obligations, because of other conflicting goals: he may lose with respect to the outcome of this contract, but may possibly win across contracts.

As for the counterparty c , we have:

$v_c(O_{b,c}(f, d)) > 0$: an obligation is an asset for the counterparty
 $v_c(f) > v_c(O_{b,c}(f, d))$: c benefits from f
 $v_c(Fulf_{b,c}(f, d)) \leq 0$: c may acquire obligations after fulfillment
 $v_c(Viol_{b,c}(f, d)) \geq 0$: c may obtain compensations after violation

Note that both fulfillments and violations may bring no value if they have no further consequences in the contract.

In a rough attempt to model the decision making process of a counterparty of an obligation whose deadline was violated, we could state that he should denounce (and thus obtain the obligation's violation) if⁵

$$v_c(f) + v_c(Fulf_{b,c}(f, d)) < v_c(Viol_{b,c}(f, d)).$$

We consider that valuations may possibly vary with time. Were that not the case, the above condition would only need to be checked right after d , at which point the counterparty would either denounce or decide to wait indefinitely for the bearer to fulfill his obligation. For instance, we believe that it makes sense to think of $v_c(f)$ as possibly decreasing with time (like a resource that should be available but is not yet). Even when the above condition does not hold, the counterparty may still prefer to tolerate the less preferred situation of failure for matters of conflicting goals (just as with the bearer).

Until now we have discussed the possibility of agents (both bearers and counterparties) deciding on breach over compliance (either by assessing intra-contract consequences or by inter-contract conflicts). But in scenarios enriched with social features agents can exploit, it may be the case that agents decide to behave cooperatively even when they have to bear a contained disadvantage. In such settings, more than being altruistic, agents may try to enhance their trust awareness in the community, from which they will benefit in future interactions or contracts.

⁵ We assume there is no cost associated with the denouncing action.

5 Implementation and Practical Issues

The logical relationships expressed above provide us a formalism to define directed deadline obligations. However, in order to monitor contracts at run-time, we need to ground this semantics into a reasoning engine capable of responding to events in a timely fashion. That is, elements describing obligation states should allow us to reason about those states *as soon as they occur*.

A natural choice we have made before [3] is the use of a rule-based inference engine, with which the following (forward-chaining) rules can be defined to implement the semantics of directed obligations with livelines and deadlines⁶:

- $O_{b,c}(f, l, d) \wedge f \wedge \neg l \rightarrow LViol_{b,c}(f, l, d)$
- $LViol_{b,c}(f, l, d) \wedge l \wedge \neg Den_{c,b}(f, l, d) \rightarrow Fulf_{b,c}(f, l, d)$
- $LViol_{b,c}(f, l, d) \wedge Den_{c,b}(f, l, d) \wedge \neg l \rightarrow Viol_{b,c}(f, l, d)$
- $O_{b,c}(f, l, d) \wedge l \wedge \neg LViol_{b,c}(f, l, d) \wedge f \wedge \neg d \rightarrow Fulf_{b,c}(f, l, d)$
- $O_{b,c}(f, l, d) \wedge d \wedge \neg f \rightarrow DViol_{b,c}(f, l, d)$
- $DViol_{b,c}(f, l, d) \wedge f \wedge \neg Den_{c,b}(f, l, d) \rightarrow Fulf_{b,c}(f, l, d)$
- $DViol_{b,c}(f, l, d) \wedge Den_{c,b}(f, l, d) \wedge \neg f \rightarrow Viol_{b,c}(f, l, d)$

With this approach, we assume an immediate assertion of facts and deadlines when they come into being. Furthermore, rules are expected to be evaluated in every working memory update (e.g. right after a fact is asserted), in order to produce the indicated conclusions, which are added to the normative state in a cumulative fashion. To detect the moment at which the *before* relation holds, we translated terms of the form $(e_1 B e_2)$ into a conjunction $e_1 \wedge \neg e_2$. The fourth rule demanded for a more careful construction, since we had two consecutive before relations – we needed to ensure that there was no liveline violation when having both l and f .

5.1 Reasoning with Time

In business contracts it is common to have deadlines that are dependent on the fulfillment date of other obligations. Therefore, instead of having fixed (absolute) dates, these may at times be relative, calculated according to other events. CISG [1] expresses this by saying that dates can be *determinable* from the contract:

Article 33: The seller must deliver the goods: (a) if a date is fixed by or determinable from the contract, on that date; (b) if a period of time is fixed by or determinable from the contract, at any time within that period [...]

Article 59: The buyer must pay the price on the date fixed by or determinable from the contract [...]

It is therefore useful to timestamp each event: facts, fulfillments and violations. For that purpose, $Fulf_{b,c}(f, l, d)^t$ will be used to indicate that b has fulfilled at time point t its obligation towards c to obtain f between l and d ; similarly for

⁶ The simpler case of directed deadline obligations is a simplification over these rules.

$Viol_{b,c}(f, l, d)^t$. Since a fact itself has now a timestamp attribute, for ease of reading we will write fact f achieved at time point t as $Fact(f)^t$. A denounce will also be written $Den_{c,b}(f, l, d)^t$.

Norms will be based on these elements and on their time references in order to prescribe other obligations with relative deadlines. For instance,

$$Fulf_{b,c}(Deliver(x, q), _, _)^t \rightarrow O_{c,b}(Pay(price), t, t + 10)$$

means that once agent b has fulfilled his obligation to deliver q units of x to agent c , the latter is obliged to pay the former within a period of 10 time units.

5.2 Re-implementing Rules

We also need to update our rules in order to stamp each generated event. In fact, having timestamps also allows us to implement such rules in a way that has a closer reading to the LTL *before* operator:

- $O_{b,c}(f, l, d) \wedge Fact(f)^t \wedge t < l \rightarrow LViol_{b,c}(f, l, d)$
- $LViol_{b,c}(f, l, d) \wedge l \wedge \neg(Den_{c,b}(f, l, d)^u \wedge u < l) \rightarrow Fulf_{b,c}(f, l, d)^l$
- $LViol_{b,c}(f, l, d) \wedge Den_{c,b}(f, l, d)^u \wedge u < l \rightarrow Viol_{b,c}(f, l, d)^u$
- $O_{b,c}(f, l, d) \wedge Fact(f)^t \wedge l < t \wedge t < d \rightarrow Fulf_{b,c}(f, l, d)^t$
- $O_{b,c}(f, l, d) \wedge d \wedge \neg(Fact(f)^t \wedge t < d) \rightarrow DViol_{b,c}(f, l, d)$
- $DViol_{b,c}(f, l, d) \wedge Fact(f)^t \wedge \neg(Den_{c,b}(f, l, d)^u \wedge u < t) \rightarrow Fulf_{b,c}(f, l, d)^t$
- $DViol_{b,c}(f, l, d) \wedge Den_{c,b}(f, l, d)^u \wedge \neg(Fact(f)^t \wedge t < u) \rightarrow Viol_{b,c}(f, l, d)^u$

This kind of approach has the benefit of relaxing the rule evaluation policy: rules do not have to be evaluated after each working memory update, since we are checking the timestamps of each event (see also [3]).

5.3 Example Contract

Considering a two-party business scenario, a contract should be beneficial for both involved parties. Therefore, both are obliged to bring about certain facts (e.g. payments or deliveries) in specific situations, and those facts should benefit the obligations' counterparties. The contract will typically specify remedies for breach situations (such as those pointed out at CISG). For the sake of illustration, we present a possible buyer-supplier contract: agent S commits to supply agent B , whenever he orders, good X for 7.5 per unit. The norms below define this particular contractual relationship. Agent S is supposed to deliver the ordered goods between 3 to 5 days after the order (norm $n1$), and agent B shall pay within 30 days (norm $n2$). Furthermore, if agent B does not pay in due time, he will incur in a penalty consisting of an obligation to pay an extra 10% on the order total (norm $n3$). Finally, if agent S violates his obligation to deliver, the contract will be canceled (norm $n4$).

- (n1) $Fact(Order(X, q))^w \rightarrow O_{S,B}(Deliver(X, q), w + 3, w + 5)$
- (n2) $Fulf_{S,B}(Deliver(X, q), l, d)^w \rightarrow O_{B,S}(Pay(q * 7.5), w, w + 30)$

- (n3) $DViol_{B,S}(Pay(p), l, d) \rightarrow O_{B,S}(Pay(p * 0.10), d, d + 30)$
 (n4) $Viols_{S,B}(Deliver(X, q), l, d)^w \rightarrow Cancel_contract$

Note that the interest applied on payments is automatic once a deadline violation is detected (norm *n3*). On the other hand, a contract cancellation (norm *n4*) requires that agent *B* denounces the inability of agent *S* to fulfill the delivery. It is therefore up to agent *B* whether to wait further and accept a delayed delivery or not. If the agreed upon contract conditions are important enough, allowing a counterparty deviation (and hence taking a cooperative attitude regarding the compliance of the contract) may be a good decision.

Different kinds of situations may be easily modeled using this kind of norms. Moreover, using flexible deadlines also ensures a degree of freedom for agents to make decisions in the execution phase of contracts, which is important for dealing with business uncertainty.

6 Related Work

Most implementations of norms in multi-agent systems ignore the need for having directed obligations from bearers to counterparties. The most likely reason for this is that in those approaches obligations are seen as (implicitly) directed from an agent to the normative system itself. It is up to the system (e.g. an electronic organization [11] or an electronic institution [12]) to detect violations and to enforce the norms which are designed into the environment (in some cases they are even regimented in such a way that violation is not possible). On the contrary, our flexible approach towards an Electronic Institution allows agents to define the norms that will regulate their mutual commitments.

Other authors have proposed different lifecycles for commitments and deontic operators. Directed social commitments are modeled in [13], in the context of dialogical frameworks. Violated commitments resort to their cancellation, which may bring sanctions. An interesting issue that is explicit in the model is the possibility for the bearer to cancel his commitment, allowing the counterparty to apply sanctions; also, updating is allowed through cancellation of the commitment and creation of a new one. A more compact model is presented in [14], also considering the possibility to update commitments. However, fulfillment and violation are not dealt with explicitly in this model; instead, a commitment is discharged when fulfilled, or else may be canceled.

Taking a cooperative approach to contract fulfillment, in [15] an obligation lifecycle model includes states that are used in a contract fulfillment protocol. Agents communicate about their intentions to comply with obligations, and in this sense an obligation can be refused or accepted. After being accepted, the obligation may be canceled or complied with. These states are obtained according to the performance of a contractual relationship. Our model should also require that agents communicate their intentions regarding an obligation with a violated deadline. In fact, CISG's Article 48 seems to go in this direction, in order to protect the bearer's efforts toward a late fulfillment of the obligation.

The need to identify two opposite roles in deontic operators is not exclusive of obligations. In [5] the concept of directed permission is described on the basis of *interference* and *counter-performance*. If a party is permitted by another to bring about some fact, the latter is not allowed to interfere with the attempt of the former to achieve that fact. The authors also sustain a relation between directed obligations and directed permissions: $O_{b,c}(f) \rightarrow P_{b,c}(f)$, that is, if an agent b has an obligation towards an agent c , then b is permitted (by c) to bring about the obliged fact and c is not permitted to interfere. This is very important in international trade transactions, especially when storage costs can be high. Some evidence from CISG [1] brings us once more the same insight:

Article 53: The buyer must pay the price for the goods and take delivery of them [...]

Article 60: The buyer's obligation to take delivery consists: (a) in doing all the acts which could reasonably be expected of him in order to enable the seller to make delivery; and (b) in taking over the goods.

In this case the permission is described in terms of an obligation of the counterparty (the buyer).

Our model of directed obligations with livelines and deadlines has some connections with research on real-time systems, where a time-value function values a task execution outcome depending on the time when it is obtained. *Soft* real-time systems use soft deadlines: obtaining the result after the deadline has a lower utility. In contrast, for *hard* real-time systems the deadline is crisp: after it, the result has no utility at all, and missing the deadline can have serious consequences. Our approach seems to be soft with a hard-deadline discretionally declared by the counterparty of the task to achieve. Deadline goals are also analyzed in [16] in the context of goal-directed and decision-theoretic planning. Goals are given a temporal extent and can be partially satisfied according to this temporal component. The authors propose a *horizon* time point somewhere after the deadline, after which there will be no benefit in achieving the goal. In our case the horizon is not static, but can be defined by the counterparty.

A model for commitment valuations, on which we have based our decision-making prospect, has been proposed in [10]. However, while their work is centered on checking correctness of contracts, we focus on valuations in the course of a contract execution. We do not assume that a contract is correct from a fairness point of view. This difference in concerns has brought divergent considerations when valuating fulfillment and violation states.

Other authors have studied agent decision-making regarding norm compliance. For instance, violation games, put in perspective of a game-theoretic approach to normative multi-agent systems in [17], model the interactions between an agent and the normative system that is responsible to detect violations and sanction them accordingly. That line of research analyses how an agent can violate obligations without being sanctioned. In our case, while we assume that temporal violations are always detected, we explore decision-making from the point of view of both the bearer and the counterparty of a directed obligation.

7 Conclusions

In cooperative B2B Virtual Organizations, contracts specify, through obligations, the interdependencies between different partners, and provide legal options to which parties can resort in case of conflict. However, when this joint activity aims at pursuing a common goal, the successful performance of business benefits all involved parties. Therefore, when developing automated monitoring tools, one should take into account that agents may be cooperative enough to allow counterparties' deviations.

Taking this into account, in this paper we have presented a novel model for contractual obligations – *directed deadline obligations*. Following a claimant theory approach, the directed aspect concerns the need to identify the agent who will be authorized to react in case of non-fulfillment. We started from previous theoretical approaches to model such authorizations, and developed a more concrete formalization by linking authorizations with a flexible model of deadlines. Obligation violations are now dependent on the counterparty motivation to claim them. We have also considered in our model smoother authorizations.

Our approach is based on real-world evidence from business contracts (namely the United Nations Convention on Contracts for the International Sale of Goods), which denotes a flexible and even cooperative facet of trade contracts. This facet extends to the concept of B2B Virtual Organizations, wherein different parties come together to share a business goal that is achievable through the cooperative fulfillment of a common contract.

We addressed the important issue of agent decision-making, which is enriched by our model of authorizations. Both parties involved in a directed deadline obligation may have a say regarding its violation. When considering obligations as interlinked through norms in a contract, agents should evaluate the consequences of fulfillment and violation states as prescribed in the contract. Furthermore, in “socially rich” environments, agents should explore the value of future relationships by enhancing their perceived trustworthiness and predisposition to facilitate compliance, something that is made possible by our directed deadline obligations approach.

Acknowledgments. The first author is supported by FCT (Fundação para a Ciência e a Tecnologia) under grant SFRH/BD/29773/2006.

References

1. UNCITRAL: United nations convention on contracts for the international sale of goods, cisg (1980)
2. Broersen, J., Dignum, F., Dignum, V., Meyer, J.J.C.: Designing a deontic logic of deadlines. In: Lomuscio, A., Nute, D. (eds.) DEON 2004. LNCS (LNAI), vol. 3065, pp. 43–56. Springer, Heidelberg (2004)
3. Lopes Cardoso, H., Oliveira, E.: A context-based institutional normative environment. In: Hübner, J.F., Boissier, O. (eds.) AAMAS'08 Workshop on Coordination, Organization, Institutions and Norms in agent systems (COIN), Estoril, Portugal, pp. 119–133 (2008)

4. Ryu, Y.U.: Relativized deontic modalities for contractual obligations in formal business communication. In: 30th Hawaii International Conference on System Sciences (HICSS), Hawaii, USA, vol. 4, pp. 485–493 (1997)
5. Tan, Y.-H., Thoen, W.: Modeling directed obligations and permissions in trade contracts. In: Proceedings of the Thirty-First Annual Hawaii International Conference on System Sciences, vol. 5. IEEE Computer Society, Los Alamitos (1998)
6. von Wright, G.: Deontic logic. *Mind* 60, 1–15 (1951)
7. Emerson, E.A.: Temporal and modal logic. In: Leeuwen, J.v. (ed.) *Handbook of Theoretical Computer Science. Formal Models and Semantics*, vol. B, pp. 995–1072. North-Holland Pub. Co./MIT Press (1990)
8. Herrestad, H., Krogh, C.: Obligations directed from bearers to counterparties. In: Proceedings of the 5th international conference on Artificial intelligence and law, College Park, Maryland, United States, pp. 210–218. ACM, New York (1995)
9. Dignum, F.: Autonomous agents with norms. *Artificial Intelligence and Law* 7(1), 69–79 (1999)
10. Desai, N., Narendra, N.C., Singh, M.P.: Checking correctness of business contracts via commitments. In: Proc. 7th Intl. Joint Conf. on Autonomous Agents and Multiagent Systems, Estoril, Portugal, IFAAMAS, pp. 787–794 (2008)
11. Vázquez-Salceda, J., Dignum, F.: Modelling electronic organizations. In: Mařík, V., Müller, J.P., Pěchouček, M. (eds.) *CEEMAS 2003. LNCS (LNAI)*, vol. 2691, pp. 584–593. Springer, Heidelberg (2003)
12. Esteva, M., Rodríguez-Aguilar, J.A., Sierra, C., Garcia, P., Arcos, J.L.: On the formal specifications of electronic institutions. In: Sierra, C., Dignum, F.P.M. (eds.) *AgentLink 2000. LNCS (LNAI)*, vol. 1991, pp. 126–147. Springer, Heidelberg (2001)
13. Pasquier, P., Flores, R.A., Chaib-Draa, B.: Modelling flexible social commitments and their enforcement. In: Gleizes, M.-P., Omicini, A., Zambonelli, F. (eds.) *ESAW 2004. LNCS (LNAI)*, vol. 3451, pp. 139–151. Springer, Heidelberg (2005)
14. Wan, F., Singh, M.P.: Formalizing and achieving multiparty agreements via commitments. In: Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems, The Netherlands, pp. 770–777. ACM Press, New York (2005)
15. Sallé, M.: Electronic contract framework for contractual agents. In: Cohen, R., Spencer, B. (eds.) *Advances in AI: 15th Conf. of the Canadian Soc. for Computational Studies of Intelligence*, pp. 349–353. Springer, Heidelberg (2002)
16. Haddawy, P., Hanks, S.: Utility models for goal-directed, decision-theoretic planners. *Computational Intelligence* 14(3), 392–429 (1998)
17. Boella, G., van der Torre, L.: A game-theoretic approach to normative multi-agent systems. In: Boella, G., van der Torre, L., Verhagen, H. (eds.) *Normative Multi-agent Systems (NorMAS07). Dagstuhl Seminar Proceedings*, vol. 07122 (2007)