Meeting the Deadline: Why, When and How

Frank Dignum, Jan Broersen, Virginia Dignum and John-Jules Meyer
Institute of Information and Computing Sciences
Utrecht University,
email: {dignum/broersen/virginia/jj}@cs.uu.nl

1 Introduction

A normative system is defined as any set of interacting agents whose behavior can usefully be regarded as norm-directed [5]. Most organizations, and more specifically institutions, fall under this definition. Interactions in these normative systems are regulated by normative templates that describe desired behavior in terms of deontic concepts (obligations, prohibitions and permissions), deadlines, violations and sanctions. Agreements between agents, and between an agent and the society, can then be specified by means of contracts. Contracts provide flexible but verifiable means to integrate society requirements and agent autonomy, and are an adequate means for the explicit specification of interactions [9]. From the society perspective, it is important that these contracts adhere to the specifications described in the model of the organization. Therefore, the specification languages must be formalized to enable formal verification.

In [8] we presented the logic LCR, based on deontic temporal logic. LCR is an expressive logic for describing interaction in multi-agent systems, including obligations with deadlines. *Deadlines* form an important type of norms in most interactions between agents. The intuitive idea of a deadline is that an agent should perform an action before a certain condition holds. The obligation to perform the action starts at the moment the deadline becomes active. E.g. when a contract is signed or approved. If the action is not performed in time a violation of the deadline occurs. It can be specified independently what measure has to be taken in this case.

In previous work, we have argued the use of a declarative specification of deadlines, as it facilitates the check of compliance and enables the reasoning about norms before the planning process determines the next sequence of actions [3]. In this paper we investigate the concept of deadline in more detail.

The paper is organized as follows. In section 2, we will discuss the basic intuitions of deadlines. Section 3 presents a first intuitive formalization for deadlines. In section 4, we look at a more complex model for deadlines trying to catch some more practical aspects. Finally, in section 5 we present issues for future work and our conclusions.

2 Basic choices for the formalization of deadlines

In this section we study some choices to make when developing a formal model for deadlines. Although the basic idea of a deadline is very simple it appears that the details are quite intricate. We suggest that one of the reasons is that in order to model deadlines, we need to model a *causal* relation between non-fulfilment of an obligation and certain violation conditions. Causal relations are notoriously hard to formalize. Figure 1 pictures the situation.

The figure shows several possible futures from a point where a deadline is in force. In some futures the required action does not take place and a violation results after the deadline is reached. For other futures, the action takes place before the deadline is reached, and no violations appear after the action.

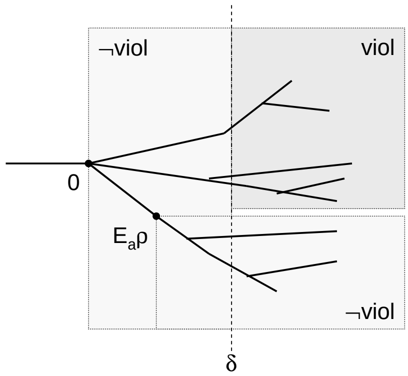

Fig. 1. *The semantics of deadlines*

In the remainder of the paper, we assume familiarity with the branching time temporal logics CTL and CTL^*. The reader can find the definitions for these branching time temporal logics in the literature (e.g. [1, 4, 2]). Here, we only give the informal interpretations of the operators. The operators E and A are path quantifiers, ranging over possible futures. The operators U (strong Until), X (neXt), G (at all points at the future), and F (at some point in the future), are linear time operators ranging over moments of individual possible futures.

In this paper we denote a deadline for agent a saying that it is obliged to achieve the condition ρ before δ holds, by the formula $O_a(\rho \leq \delta)$. We will give a formal definition of the semantics of this formula after, in the next sections, we have discussed some basic choices to make.

2.1 Do obligations persist after the deadline?

A first distinction we make is between deadline obligations that are discharged by a failure to meet the deadline, and deadline obligations where the obligation is not discharged at the deadline. For a deadline of the first type it makes no sense to perform the action after the deadline passes. E.g., submitting a paper after the deadline of a conference has no effect. An example of the second type is the situation where one has to pay mortgage before the end of the month. Also when no payment has been made, this obligation persists until after the end of the month. We see this second type of deontic deadlines as the conjunction of a standard obligation and a deadline obligation of the other type. In this paper we only consider deadline obligations of the first type, i.e., deadlines where the obligation is no longer in force after the deadline. The second option can be modelled indirectly by adding a rule that states that when a deadline is violated a new obligation arises to perform the action immediately.

2.2 What if the deadline is never or immediately met?

We first consider the case where δ equals $\perp$. Clearly, $\perp$ is a condition that will be never met. Then, a natural question is, whether it is actually possible to have a deadline obligation for a deadline that never occurs. One could choose to say that this is impossible, which leads to the desired property (1) $\models \neg O_a(\rho \leq \perp)$. Another possibility is to say that such an obligation is possible, but without any force, and thus void. This corresponds to the property (2) $\models O_a(\rho \leq \perp)$. This obligation is void, because it can never be violated. A similar situation occurs in standard deontic logic, where we have $\models O\top$, which corresponds with the void obligation for a tautology (also something that can never be violated). Our formalization in section 3 satisfies this property. In section 4 we add the causal relation that if an obligation is complied to, no violations of that obligations can occur afterwards. Then, if we choose (as in section 3) that we comply to $O_a(\rho \leq \bot)$ trivially (it is void), than the causal approach dictates that we obey (3) $\models O_a(\rho \leq \bot) \rightarrow AG\neg Viol$.

One could argue that an obligation of the form $O_a(\rho \leq \bot)$ should not be void at all. For some agents, the impossibility of the deadline condition might not imply that it is free to refrain from ρ forever (even if a violation cannot occur in finite time). This leads to the view expressed by (4) $\models O_a(\rho \leq \bot) \rightarrow AF\rho$, saying that if we have an obligation for a deadline that never occurs, the agent will have to comply anyway at some future point in time (where that point can of course be arbitrarily far in the future).

Now consider the case where δ equals $\top$. This means that the deadline condition is met in the current state. Then, the only possibility to comply to the obligation, is that ρ is already achieved. This leads to the desired property (5) $\models O_a(\rho \leq \top) \rightarrow Viol(a) \vee \rho$.

To be complete, we can also consider what happens if the deadline already held in the past. We do not think it is very useful to have a valid obligation in this case, because it can never be fulfilled unless it was fulfilled already at that time. So, in our definition of deadlines we assume the deadline to be a formula that should hold at some point in the future.

2.3 What if the accomplishment is never or trivially achieved?

First of all we have to determine whether an obligation to achieve a state ρ is fulfilled whenever the state ρ holds. It is possible that the state ρ holds accidentally without any effort or intention of the agent for whom the obligation holds. E.g. if a person is obliged to write the introduction of a paper, fails to do so, but a co-author writes the introduction (because she is tired of waiting for that person). Did the person fulfill his obligation or not? We believe that if the obligation is *directed* to an agent, then that agent should *see to it that* the obligation is fulfilled. We therefore introduce a modal operator $E_a\rho$ that links the establishment of the required state to the person to which the obligation is directed. We base our use of the *stit* operator E_a on the work of Pörn [6], but add the property that agents cannot achieve tautologies and the state is achieved one step after the action is initiated ($E_a\rho \rightarrow X\rho$)!

Our next question concerns the case where the achievement can never be reached. This is possible if the state to achieve is not under control of the agent. A special case is when that state is $\bot$. Our first option is again to state that obligations of the form $O_a(\bot \leq \delta)$ are impossible or inconsistent. After all, it seems reasonable to take the position that one can never be obliged to achieve the impossible. This leads to the desired property (6) $\models \neg O_a(\bot \leq \delta)$, which is similar to standard deontic logic's D-axiom $\neg O\bot$. However, we might also take the position that one can have an obligation to achieve the impossible. But, since $O_a(\bot \leq \delta)$ expresses that we have to achieve the impossible before the deadline condition δ occurs, we have to conclude that this leads to the view that there will certainly be a violation whenever δ occurs for the first time. This leads to the property: (7) $\models O_a(\bot \leq \delta) \rightarrow \neg E(\neg\delta U(\delta \wedge \neg Viol))$.

Finally we consider the case where the accomplishment is $\top$. Because $\top$ cannot be achieved in our formalism (it is already true), this deadline is not valid.

3 Formalizing deadlines

After having discussed some major choices for modelling deadlines in the previous sections we will present a first intuitive logical formalization.

The deontic aspect of deadlines is formalized by introducing a set $\mathcal{A}$ of agent identifiers and a propositional constant $Viol(a)$ for each $a \in \mathcal{A}$ in the language of CTL. The general idea is that the violation condition holds (i.e., the propositional constant $Viol(a)$ is true) at those moments where agent a violates a deontic deadline. This enables us to reason about violations explicitly, and about what to do if they occur, which is a distinctive feature of deontic reasoning. We model deadlines conditions as propositions. This seems a reasonable choice given that we do not want to model a deadline in a logic of explicit time (real time). Our view is more abstract, and a deadline is simply a condition true at some point in time. We use the symbols δ and γ to denote deadline propositions.

$E_a\rho$ indicates that the agent a sees to it that ρ becomes true. If $E_a\rho$ is true at a point in time than ρ will be true at the next point in time. We use the symbols ρ and σ for propositions that embody some kind of accomplishment being established before a deadline condition occurs.

Given a CTL model $M = (S, \mathcal{R}, \pi)$, with S a non-empty set of states, $\mathcal{R}$ an accessibility relation such that the reflexive transitive closure of $\mathcal{R}$ is a total tree relation, and π a valuation of propositional atoms, a path t in M is a sequence $t = s_0, s_1, s_2, \ldots$ such that for every $i \geq 0$, s_i is an element of S and $s_i \mathcal{R} s_{i+1}$, and if t is finite with s_n its final situation, then there is no situation s_{n+1} in S such that $s_n \mathcal{R} s_{n+1}$. We say that a path starts at s iff $s_0 = s$. If $t = s_0, s_1, s_2, \ldots$ is a path in M, then we denote s_i by $t(i)$ $(i \geq 0)$. We denote the tail of the path t starting at state $t(i)$ by $t[i]$. Let M be a CTL model, s a state, and t a full path. A straightforward modal semantics for the operator $O_a(\rho \leq \delta)$ is then defined as follows:

$$\begin{aligned} M, s \models O_a(\rho \leq \delta) \Leftrightarrow\ & \forall t \text{ with } t(0) = s, \\ & \forall j: \text{ if } M, t(j) \models \delta : \\ & \ \text{and } \forall 0 \leq i < j : M, t(i) \models \neg E_a\rho, \\ & \ \text{then } M, t(j) \models Viol(a) \end{aligned}$$

This says: if at some future point the deadline occurs, and until then the result has not yet been achieved, then we have a violation at that point. This semantic definition is equivalent to the following definition as a reduction to CTL:

$$O_a(\rho \leq \delta) \equiv_{def} \neg E(\neg E_a\rho U(\delta \wedge \neg Viol(a)))$$

This formula just expresses the negation of the situation that should be excluded when a deontic deadline is in force. In natural language this *negative* situation is: 'δ becomes true at a certain point, the achievement has not been met until then, and there is *no* violation at δ'. This shows that it is fairly easy to show the equivalence of the semantic definition and the definition in terms of CTL (details left to the reader). The above defined deadline operator satisfies properties 2 and 5 discussed in the previous section (after some non-essential adjustments concerning the operator $E_a\rho$).

However, despite the nice properties and the simple and elegant representation of the concepts, the definition does not cover the intuitions of figure 1 completely. This becomes apparent when we look at a situation in which an agent a achieves ρ before a certain condition δ becomes true. Whenever this appears to be true it follows that a has the obligation to achieve ρ. I.e., the fact that an agent will achieve something implies that he is obliged to achieve it. One can easily check that this follows from the definition.

When analyzing this problem one sees that the causal link that we intuitively suppose to be present between the non-achievement (in time) and the violation state and the achievement and the non-violation state is not captured by the single implication of the semantic definition. In the next section we will propose an extended definition that tries to establish this causal link between the achievement and the violation state.

4 The causal approach

In [8] we already tried to capture some of the causal link between the appearance of the achievement and the violations. However that formalization did not force the fact that there can never be a violation of the obligation before the deadline condition holds. It also allows situations where ρ is achieved while there is still a violation after the deadline condition. Somehow we have to "close" the possible worlds in a way that either we have the achievement and no violation after that or a violation and no achievement before the deadline. In this way we approach most closely that the achievement of ρ *causes* the $\neg Viol(a)$.

$$
\begin{aligned}
M, s \vDash O_a(\rho < \delta) \textbf{ iff } & \forall t \text{ with } t(0) = s : \exists j : \\
& M, t(j) \models \delta \textbf{ and} \\
& (\forall 0 \leq k \leq j : M, t(k) \vDash \neg Viol(a)) \textbf{ and} \\
& ((\exists 0 \leq k < j : M, t(k) \vDash E_a\rho \textbf{ and } M, t(k) \vDash AG\neg Viol(a)) \textbf{ or} \\
& (\forall 0 \leq k < j : M, t(k) \nvDash E_a\rho \textbf{ and } M, t(j) \vDash AGViol(a)))
\end{aligned}
$$

The first part of the definition is equal to the first definition given in section 3. However, we now require $\neg Viol(a)$ to hold until the deadline condition δ is reached. Finally, we have a disjunction that should be read like an "if-then-else". Either $E_a\rho$ holds (in time) and $\neg Viol(a)$ holds from that point on or $E_a\rho$ does not hold before δ and $Viol(a)$ is true after the deadline. Note that the or is exclusive, because either ρ is achieved or not, but not both can be true at the same time.

We can express the above semantic definition in terms of a formula as well:

$$
\begin{aligned}
O_a(\rho < \delta) \equiv_{def} A(& \\
& (\neg Viol(a) \wedge \neg\delta) U \delta) \wedge \\
& ((((\neg E_a\rho \wedge \neg\delta) U E_a\rho) \wedge (\neg E_a\rho U (E_a\rho \wedge AG\neg Viol(a))) \vee \\
& (((\neg E_a\rho \wedge \neg\delta) U \delta) \wedge (\neg\delta U (\delta \wedge AGViol(a))))))
\end{aligned}
$$

The formula states that a is obliged to achieve ρ before the deadline δ can be defined to hold whenever the deadline condition is reached in the future then $Viol(a)$ does not hold until that time and either ρ is achieved before δ and after the first time that ρ is achieved $\neg Viol(a)$ holds in all possible futures from that point on or ρ is not achieved before δ and $Viol(a)$ holds in every possible future after the deadline condition becomes true.

Although the final formula looks quite different from the one given in section 3.3 it is not very difficult to prove that it is actually stronger than that one and thus the obligation as defined above implies the one defined in section 3.3.

Notice that using the above definition the deadline obligation does not (need to) hold anymore after the required state is achieved. However, we keep the information of having fulfilled the obligation through the fact that in every possible future state $\neg Viol(a)$ holds.

5 Practical aspects of deadlines

In this section we will very briefly mention a few aspects that start playing a role when looking at more concrete aspects of deadlines.

The first aspect is the violation predicate. In this paper the $Viol$ predicate has only one parameter, the agent a. However, we would actually like to tie the violation to a specific obligation incurred at a specific moment in time. This is necessary to distinguish two obligations for the same agent that might only differ in the timing. E.g. the obligation to pay the rent before the end of the month occurs every month. But each month it is a different obligation.

Closely related to the above item is the point that we made violations (and non-violations) persistent over time. Once a deadline is violated, this violation will never disappear again. This seems a bit contradictory to common practice where sanctions are defined as obligations, conditional on the occurrence of a violation, in order to make it possible for violations to be redeemed. So, we make a difference between a violation that has not been "made up for" yet and one for which a sanction has been exercised already.

A second item that is important in practice is that obligations are often conditional and/or repeated. The above example on paying the rent is a very typical case of a repeated obligation. The whole obligation to pay rent, however, can be made conditional on the fact that the house is properly maintained by the owner. Related to this aspect is that more temporal conditions can be specified for the achievement. E.g. the salary should be paid between the 25th and the end of each month.

Although we represent the deadline condition as a proposition in this paper, often it contains a relative temporal expression such as "the book should be paid within one week after delivery". In order to express this type of conditions one should have a more powerful language in which explicit reference to time can be made.

A last item to mention here is the use of discrete time in our model. This is particularly important to decide on the exact moment when a violation arises. In a model with continuous time the achievement of a fact (an action) has to have a duration (whereas the achievement in our model is always in one time step). So the definition of $E_a\rho$ has to be changed. In the other hand we can in this model with continuous time determine a violation before the deadline if it is impossible to achieve the required state before the deadline condition anymore.

6 Conclusions

In this paper we have shown that the use of a violation predicate is in principle enough powerful to account for the deontic aspect of the deadlines. Of course a temporal logic is needed to account for the temporal aspects. Finally we used the *stit* operator E_a to relate the achievement of a state to an agent. This is important, because we consider the deadlines to be directed towards an agent and thus this agent has the responsibility to fulfill it. We do not use dynamic logic to model explicit actions in order to keep the model as abstract as possible. However, an obvious connection between the operator presented and dynamic logic can be made through the use of Segerberg's *b*ringing it about operator [7].

We have also shown that a correct definition of deadlines in the formalism requires a modelling of the intuitive causal relation between the occurrence of the action before the deadline and the violation state. This causal relation makes the formal definition of a deadline quite complicated, although the simple intuitive picture of the semantics (given in section 2) is still valid.

References

1. E.M. Clarke, E.A. Emerson, and A.P. Sistla. Automatic verification of finite-state concurrent systems using temporal logic specifications. *ACM Transactions on Programming Languages and Systems*, 8(2), 1986.
2. E.M. Clarke, O. Grumberg, and D. Long. Verification tools for finite-state concurrent systems. In *A decade of concurrency*, volume 803 of *Lecture Notes in Computer Science*, pages 124–175. Springer, 1993.
3. F. Dignum and R. Kuiper. Obligations and dense time for specifying deadlines. In *Proceedings of thirty-First HICSS, Hawaii*, 1998.
4. E.A. Emerson. Temporal and modal logic. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science, volume B: Formal Models and Semantics*, chapter 14, pages 996–1072. Elsevier Science, 1990.
5. A.J.I. Jones and M. Sergot. On the characterization of law and computer systems: The normative sytems perspective. In J.-J.Ch. Meyer and R.J. Wieringa, editors, *Deontic Logic in Computer Science: Normative System Specification*, pages 275–307. John Wiley and Sons, 1993.
6. I. Pörn. *Action Theory and Social Science: Some Formal Models*. D. Reidel Publishing Company, 1977.
7. K. Segerberg. Bringing it about. *Journal of Philosophical Logic*, 18(4):327–347, 1989.
8. F. Dignum H. Weigand V. Dignum, J.J. Meyer. Formal specification of interaction in agent societies. In *2nd Goddard Workshop on Formal Approaches to Agent-Based Systems (FAABS), Maryland*, 2002.
9. E. Verharen, H. Weigand, and F. Dignum. Dynamic business models as a basis for interoperable transaction design. *Information Systems*, 22(2).