# 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 [9]. 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 [14]. From the society perspective, it is important that these contracts adhere to the specifications described in the model of the organization. If we want to automate such verifications, we have to formalize the languages used for contracts and for the specification of organizations.

In [13] we presented the logic LCR, which is based on deontic temporal logic. LCR is an expressive language for describing interaction in multi-agent systems, including obligations with deadlines. *Deadlines* are important norms in most interactions between agents. Intuitively, a deadline states that an agent should perform an action before a certain point in time. 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 advocated the use of declarative deadline specifications, as it facilitates the check for compliance to a deadline and enables reasoning about norms before the planning process determines the next sequence of actions [5]. In this paper we investigate the deadline concept in more detail.

The paper is organized as follows. Section 2 defines the variant of CTL we use. In section 3, we discuss the basic intuitions of deadlines. Section 4 presents a first intuitive formalization for deadlines. In section 5, we look at a more complex model for deadlines trying to catch some more practical aspects. Finally, in section 6 we present issues for future work and our conclusions.

## 2 Preliminaries: CTL

The reader can find the definitions for the branching time logic CTL in the literature (e.g. [3, 7, 4]). But, since we need a specific variant of the until operator, we define CTL here explicitly.

Well-formed formulas of the temporal language $\mathcal{L}_{\text{CTL}}$ are defined by:

$$\begin{aligned}
\varphi, \psi, \ldots &:= p \mid \neg\varphi \mid \varphi \wedge \psi \mid E\alpha \mid A\alpha \\
\alpha, \beta, \ldots &:= \varphi U^e \psi \mid X\varphi
\end{aligned}$$

where $\varphi, \psi$ represent arbitrary well-formed formulas, and where the $p$ are elements from an infinite set of propositional symbols $\mathcal{P}$. Formulas $\alpha, \beta, \ldots$ are called 'path formulas'. We use the superscript 'e' for the until operator to denote that this is the version of 'the until' where $\varphi$ is not required to hold for the point where $\psi$, i.e., the point where $\phi$ is excluded. However, the present state is not excluded, which means that our until operator is reflexive. This gives us the following informal meanings of the until operator:

$E(\varphi U^e \psi)$ : there is a future for which eventually, at some point $m$, the condition $\psi$ holds, while $\varphi$ holds from now until the moment before $m$

We define all other CTL-operators as abbreviations. Although we do not use all of the LTL operators $X$, $F$, and $G$ in this paper, we give their abbreviations (in combination with the path quantifiers $E$ and $A$) in terms of the defined operators for the sake of completeness. We also assume the standard propositional abbreviations.

$$\begin{array}{ll}
EF\varphi \equiv_{def} E(\top U^e \varphi) & AG\varphi \equiv_{def} \neg EF\neg\varphi \\
AF\varphi \equiv_{def} A(\top U^e \varphi) & EG\varphi \equiv_{def} \neg AF\neg\varphi \\
A(\varphi U \psi) \equiv_{def} A(\varphi U^e (\varphi \wedge \psi)) & E(\varphi U \psi) \equiv_{def} E(\varphi U^e (\varphi \wedge \psi))
\end{array}$$

The informal meanings of the formulas with a universal path quantifier are as follows (the informal meanings for the versions with an existential path quantifier follow trivially):

$A(\varphi U \psi)$ : for all futures, eventually, at some point the condition $\psi$ will hold, while $\varphi$ holds from now until then
$AX\varphi$ : at any next moment $\varphi$ will hold
$AF\varphi$ : for all futures, eventually $\varphi$ will hold
$AG\varphi$ : for all possible futures $\varphi$ holds globally

A CTL model $M = (S, \mathcal{R}, \pi)$, consists of a non-empty set $S$ of states, an accessibility relation $\mathcal{R}$, and an interpretation function $\pi$ for propositional atoms. A full path $\sigma$ in $M$ is a sequence $\sigma = s_0, s_1, s_2, \ldots$ such that for every $i \geq 0$, $s_i$ is an element of $S$ and $s_i \mathcal{R} s_{i+1}$, and if $\sigma$ 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 the full path $\sigma$ starts at $s$ if and only if $s_0 = s$. We denote the state $s_i$ of a full path

$\sigma = s_0, s_1, s_2, \ldots$ in $M$ by $\sigma_i$. Validity $M, s \models \varphi$, of a CTL-formula $\varphi$ in a world $s$ of a model $M = (S, \mathcal{R}, \pi)$ is defined as:

$$
\begin{array}{ll}
M, s \models p & \Leftrightarrow s \in \pi(p) \\
M, s \models \neg\varphi & \Leftrightarrow \text{not } M, s \models \varphi \\
M, s \models \varphi \wedge \psi & \Leftrightarrow M, s \models \varphi \text{ and } M, s \models \psi \\
M, s \models E\alpha & \Leftrightarrow \exists\sigma \text{ in } M \text{ such that } \sigma_0 = s \text{ and } M, \sigma, s \models \alpha \\
M, s \models A\alpha & \Leftrightarrow \forall\sigma \text{ in } M \text{ such that } \sigma_0 = s \text{ it holds that } M, \sigma, s \models \alpha \\
M, \sigma, s \models X\varphi & \Leftrightarrow M, \sigma_1 \models \varphi \\
M, \sigma, s \models \varphi U^e \psi & \Leftrightarrow \exists n > 0 \text{ such that} \\
 & \quad (1)\ M, \sigma_n \models \psi \text{ and} \\
 & \quad (2)\ \forall i \text{ with } 0 \leq i < n \text{ it holds that } M, \sigma_i \models \varphi
\end{array}
$$

Validity on a CTL model $M$ is defined as validity in all states of the model. If $\varphi$ is valid on a CTL model $M$, we say that $M$ is a model for $\varphi$. General validity of a formula $\varphi$ is defined as validity on all CTL models. The logic CTL is the set of all general validities of $\mathcal{L}_{\text{CTL}}$ over the class of CTL models.

## 3 Basic choices for the formalization of deadlines

In this section we study some choices to make when developing a formal model for deadlines. 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 $\mathcal{L}_{\text{CTL}}$. 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 deadline 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 $\delta$ and $\gamma$ to denote deadline propositions.

Although the basic idea of a deadline is very simple it appears that the details are 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, so called, '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 does take place before the deadline is reached, and no violations appear after the action.

We denote a deadline for agent $a$ saying that it is obliged to achieve the condition $\rho$ before $\delta$ 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.

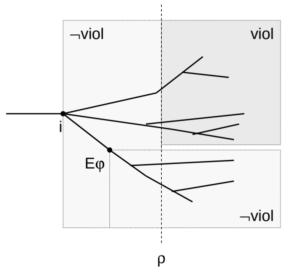


**Fig. 1.** *The semantics of deadlines*

### 3.1 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 a fine for some traffic offense by the end of the month. Also when one does not pay, the obligation to pay persists (see also the work of Brown on 'standing obligations' [2]). Yet another category are the 'repetitive obligations', where the same deadline obligation is repeated over a period of time. For example monthly mortgage payments.

### 3.2 What if the deadline is never or immediately met?

We first consider the case where $\delta$ equals $\perp$. Clearly, $\perp$ is a condition that will be never met. 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 optional property (1) $\models \neg O_a(\rho \leq \perp)$. This is the case for our deadline definition is section 5, because, in the definition given there, we assume that a deadline obligation can only be in force if the deadline condition actually occurs at some point in the future. Another possibility is to say that for any condition $\rho$ such an obligation is actually always valid, but void, i.e, without any 'force'. This corresponds to the property (2) $\models O_a(\rho \leq \perp)$. Such obligations can be considered void, because they cannot be violated; since the deadline never occurs, there will never be a point in time where non-compliance is evaluated. It might be argued that a similar situation occurs in standard deontic logic [15], 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 4 satisfies this property.

Obviously, the third possibility is that neither property (1), nor property (2) is satisfied. For instance, one could argue that an obligation for a deadline that never occurs, i.e., $O_a(\rho \leq \perp)$, is not void, but should be interpreted as follows: the impossibility of the deadline condition means that the deadline is ill-defined, but this does not imply that the agent is free to postpone his duty forever: he has to comply at some future point anyway (where that point can be arbitrarily far in the future). The corresponding formula is: (3) $\models O_a(\rho \leq \bot) \rightarrow AF\rho$.

Now consider the case where $\delta$ equals $\top$. This means that the deadline condition is met trivially, in the current state. One possible view is that in this case, we can still comply to the obligation by ensuring that also $\rho$ is met in the current state. The corresponding property is: (4) $\models O_a(\rho \leq \top) \rightarrow Viol(a) \vee \rho$.

Alternatively, we might argue that it is impossible to comply to a deadline for which the deadline condition is true *now*. For an agent, it takes some time to decide whether or not to comply, and to bring about the condition $\rho$ the obligation is concerned with. Then, if the deadline condition is true now, there is no time left for this process, and the agent will inevitably violate the obligation. In our definitions of section 4 and 5, we take this aspect into account. The corresponding property is (5) $\models O_a(\rho \leq \top) \rightarrow Viol(a)$, which is satisfied by the deontic deadline definition in sections 4 and 5. Note that under this view, the violation is not avoided if accidentally the condition $\rho$ is true in the present state. This is because under this view, conditions are linked to agents that bring them about, which is a decision they make in the previous state, as we explain later on.

Finally one short comment about the thought that we have to account for the situation that a deadline condition might have been true in the past. Clearly we do not have to consider this situation, because it is impossible to have an obligation to do something *before* something that occurred in the past.

### 3.3 What if the accomplishment is accidentally, never or trivially achieved?

First we address the question whether it counts as compliance to a deadline obligation when the condition that is obliged occurs 'accidentally'. It is possible that the state $\rho$ occurs 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 he is tired of waiting for that person). Did the person fulfill his obligation or not? If obligations are personal, should it not be the case that also the achievements $\rho$ are personal? After all, we do not want that if another agent, or 'nature', brings about the achievement, the agent with the obligation has complied. We encounter a basic choice to make here. If we do not want our obligations to be personal, we do not have to personalize the achievements. But, if we do want our obligations to be personal, we somehow have to link achievements to agents. There is a vast amount of literature about personalizing the achievement of conditions [10, 1, 8, 6]. Usually, such theories are called 'logics of action and/or agency'. Inspired by the work of Pörn [10], we use the *stit* operator $E_a\rho$, to denote that agent $a$ achieves condition $\rho$. A difference with the *stit* operator of Pörn is that in our temporal setting, performing a 'seeing to it' action takes one time-step. That is, our stit-operator obeys $\models E_a\rho \rightarrow X\rho$, and not $\models E_a\rho \rightarrow \rho$, which holds for most other agency operators.

Our next question concerns the case where the achievement can never be reached. For instance, one might think of a personal obligation for a condition not under control of an agent. An example is the condition $\bot$. Again, a first option is to say 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 optional property (6) $\models \neg O_a(\bot \leq \delta)$, which is similar to standard deontic logic's D-axiom $\neg O\bot$ [15]. 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 $\delta$ occurs, we have to conclude that this leads to the view that there will certainly be a violation whenever $\delta$ occurs for the first time. This leads to the optional property: (7) $\models O_a(\bot \leq \delta) \rightarrow \neg E(\neg\delta U^e(\delta \wedge \neg Viol(a)))$.

Finally we consider the case where the accomplishment is $\top$. How to deal with this situation depends on whether we consider the obligation to be personal or not. As discussed, for the personal case, we have to use an agency operator. In most logics of agency, $\top$ cannot be achieved by any agent ($\models \neg E_a\top$). This motivates the optional property (8) $\models \neg O_a(\top \leq \delta)$. However, if obligations are not personal, this is not necessarily intuitive. At this point we might not want to digress from standard deontic logic, where the obligation for a tautology is always valid. Thus we have the optional property (9) $\models O_a(\top \leq \delta)$.

## 4 A simple formalization

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

As mentioned, $E_a\rho$ indicates that the agent $a$ sees to it that $\rho$ becomes true. If $E_a\rho$ is true at some point in time, then $\rho$ is true at the next point in time. We use the symbols $\rho$ and $\sigma$ for propositions that embody some kind of accomplishment being established before a deadline condition occurs.

Let $M$ be a $CTL$ model, $s$ a state, and $\sigma = \sigma_0, \sigma_1, \sigma_2, \ldots$ a full path in $M$. A straightforward modal semantics for the operator $O_a(\rho \leq \delta)$ is then defined as follows:

$$\begin{array}{rl} M, s \models O_a(\rho \leq \delta) \Leftrightarrow & \forall\sigma \text{ with } \sigma_0 = s, \forall j: \\ & \text{if } M, \sigma_j \models \delta \\ & \text{and } \forall i \text{ with } 0 \leq i < j : M, \sigma_i \models \neg E_a\rho, \\ & \text{then } M, \sigma_j \models Viol(a) \end{array}$$

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 $CTL$:

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

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: '$\delta$ becomes true at a certain point, the achievement has not been met until then, and there is *no* violation at $\delta$'. 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 persists after reaching the deadline, and satisfies properties 2, 5, and 7 discussed in the previous section.

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 $\rho$ before a certain condition $\delta$ becomes true. Whenever this appears to be true it follows that $a$ has the obligation to achieve $\rho$. I.e., the fact that an agent will achieve something implies that he is obliged to achieve it.

We suggest that the source of this problem might be that we have failed to formalize the 'causal link' that intuitively relates failures to comply to the obligation and occurrences of the violation condition. In the truth condition above, we have only dealt with one direction of the implicative relation between non-compliance and violation: we have captured that when there is non-compliance, there is also a violation. But we have failed to capture a reverse implicative direction saying that only if there is non-compliance there can be violations.

In the next section we will propose an extended definition that tries to establish this causal link between non-achievements and violations.

## 5 The causal approach

In [13] we have already attempted to capture some aspects of the causal link between non-achievement and violations. However that formalization did not force the condition that there can never be a violation of the obligation before the deadline condition holds. It also allows situations where $\rho$ 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 $\rho$ *causes* the $\neg Viol(a)$.

The definition given below differs from the one in section 4 on three important points. First of all, for a deadline obligation to be valid, it now requires that the deadline condition actually occurs at some point in the future. A second crucial difference is that we strengthen the 'if' construction in the truth condition to an 'if-and-only-if' condition, by which we attempt to capture the causal relation between non-compliance and violation. This 'if-and-only-if' condition takes the form of a disjunction (the 'or' in the truth condition below) saying that either $E_a\rho$ holds (in time), meaning that there is compliance, or $E_a\rho$ does not hold before $\delta$, in which case there is non-compliance. Note that the disjunction is exclusive, because either $\rho$ is achieved or not, but not both. Finally, we require violations to persist ones they have occurred, and we require non-violations to persist when the achievement is accomplished in time, or if no deadline or achievement condition has yet occurred.

$$
\begin{aligned}
M, s \vDash O_a(\rho \leq \delta) \textbf{ iff } & \forall \sigma \text{ with } \sigma_0 = s : \exists j > 0 : \\
& M, \sigma_j \models \delta \textbf{ and } \forall 0 \leq k < j : M, \sigma_k \vDash \neg Viol(a) \wedge \neg \delta \textbf{ and} \\
& (\exists 0 \leq k < j : M, \sigma_k \vDash E_a \rho \wedge AG \neg Viol(a) \textbf{ or} \\
& (\forall 0 \leq k < j : M, \sigma_k \vDash \neg E_a \rho \textbf{ and } M, \sigma_j \vDash AG\ Viol(a)))
\end{aligned}
$$

We can express this semantic definition in terms of a CTL formula as well:

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

The lines of the formula correspond to the lines of the truth condition. The second line expresses that $\delta$ becomes true at a specific point in the future, that we consider the first time this happens, and that there cannot be a violation of the obligation until then. The third line expresses one side of the exclusive disjunction, saying that $E_a \rho$ occurs before the first $\delta$, and that there cannot be a violation afterwards. The fourth line expresses the other side of the disjunction, saying that $E_a \rho$ has not occurred before the first $\delta$, and that starting from the point where $\delta$, violations persist forever. The latter condition expresses that the information that the obligation is violated, is preserved.

In the above definition, the obligation is always discharged by the occurrence of a deadline condition. So, for this variant, the obligation does not persist until after the deadline. Furthermore, the definition obeys the properties 1, 5 and 7 of section 3.

## 6 Practical aspects of deadlines

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

The first aspect is the violation constant. In this paper the $Viol$ constant 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. This can be achieved through the addition of a unique identifier for each obligation. This definition provides a very operational means to deal with violations, as it gives explicit information about what has caused the violation and can therefore enable to reason about what are the consequences and sanctions related to the violation.

However, at the same time this unique identifier would eliminate any logical relations between obligations that are connected. E.g. someone might have an obligation to pay a conference fee while (due to budget restrictions that became clear only later) it is from now on prohibited to pay for any conference. The two norms relate to the same person and have opposite effects on the action of paying. However, if each would be modelled with a violation constant with a different identifier they could not be related and the intuitive contradiction between the two would not exist.

As a solution to this problem we could introduce violations that have the same parameters as the obligations to which they are linked. In this way it becomes possible to specify logical relations between violations of which the actor, the deadlines and the situation to be achieved are related. However, this has as consequence that the violations are now also modal operators!
A second point that comes up right away is which logical relations should hold between the violations? Do we have

$$(V_a(\rho < \delta) \wedge (\rho' \rightarrow \rho)) \longrightarrow V_a(\rho' < \delta)$$

and/or

$$(V_a(\rho < \delta) \wedge (\delta' \rightarrow \delta)) \longrightarrow V_a(\rho < \delta')$$

Of course these properties are directly coupled to the properties that we would like to have for the obligation operator. A complete investigation into this issue warrants a separate paper and therefore will not be pursued here. However we would like to point to [11] for some related work in this area.

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. This aspect could be modelled by not having the violation persistent, but have an axiom that triggers a sanction (obligation) whenever a violation occurs.

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. On 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.

## 7 Conclusions

In this paper we have shown that the use of a violation constant 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 [12].

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. N. Belnap and M. Perloff. Seeing to it that: A canonical form for agentives. *Theoria*, 54:175–199, 1988.
2. M.A. Brown. Doing as we ought: towards a logic of simply dischargeable obligations. In M.A. Brown and J. Carmo, editors, *Deontic logic, agency, and normative systems. Proceedings DEON '96*, Workshops in Computing, pages 47–65. Springer, 1996.
3. 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.
4. 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.
5. F. Dignum and R. Kuiper. Obligations and dense time for specifying deadlines. In *Proceedings of thirty-First HICSS, Hawaii*, 1998.
6. D. Elgesem. The modal logic of agency. *Nordic Journal of Philosophical Logic*, 2(2):1–46, 1997.
7. 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.
8. R. Hilpinen. On action and agency. In E. Ejerhed and S. Lindström, editors, *Logic, Action and Cognition: Essays in Philosophical Logic*, pages 3–27. Kluwer Academic Publishers, 1997.

9. 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.
10. I. Pörn. *Action Theory and Social Science: Some Formal Models*. D. Reidel Publishing Company, 1977.
11. J.-J.Ch. Meyer F. Dignum R. Wieringa, H. Weigand. The inheritance of dynamic and deontic integrity constraints. *Anals of Mathematics and Artificial Intelligence*, 3.
12. K. Segerberg. Bringing it about. *Journal of Philosophical Logic*, 18(4):327–347, 1989.
13. 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.
14. E. Verharen, H. Weigand, and F. Dignum. Dynamic business models as a basis for interoperable transaction design. *Information Systems*, 22(2).
15. G.H. von Wright. Deontic logic. *Mind*, 60:1–15, 1951.