Thomas Ågotnes
Wiebe van der Hoek
Juan A. Rodríguez-Aguilar
Carles Sierra
Michael Wooldridge

Multi-Modal CTL: Completeness, Complexity, and an Application

Abstract. We define a multi-modal version of Computation Tree Logic (CTL) by extending the language with path quantifiers E^δ and A^δ where δ denotes one of finitely many dimensions, interpreted over Kripke structures with one total relation for each dimension. As expected, the logic is axiomatised by taking a copy of a CTL axiomatisation for each dimension. Completeness is proved by employing the completeness result for CTL to obtain a model along each dimension in turn. We also show that the logic is decidable and that its satisfiability problem is no harder than the corresponding problem for CTL. We then demonstrate how *Normative Systems* can be conceived as a natural interpretation of such a multi-dimensional CTL logic.

Keywords: Computation Tree Logic (CTL), Normative Systems, Social Laws.

1. Introduction

Computation Tree Logic (CTL) is one of the most popular and successful logics in computer science [8]. CTL has been very widely applied, and has received particular prominence through the development of efficient and industrially applicable CTL model checking systems such as SMV [6].

CTL is a branching time temporal logic, and temporal operators in CTL are made by combining a *path quantifier* with a *tense modality.* The possible path quantifiers are E ("for some path"), and A ("for all paths") while the possible tense modalities are $\Diamond$ ("eventually"), $\Box$ ("always"), $\bigcirc$ ("next"), and $\mathcal{U}$ ("until"). Thus, a formula such as $\mathsf{A}\Box\phi$ expresses the fact that ϕ is an invariant, i.e., ϕ is true at every state along every future path. CTL formulae are interpreted in a state in a Kripke structure, with the accessibility relation taking the role of a next-state relation. The relation is usually required to be total, and a state can have more than one possible next state, modelling branching time.

In this paper we generalise CTL to a finite set of *dimensions* Δ. Syntactically, we have one instantiation $\mathsf{E}^\delta, \mathsf{A}^\delta$ of the path quantifiers for each

Presented by **Jacek Malinowski**; *Received* March 9, 2009

Studia Logica (2009) 92: 1–26
DOI: 10.1007/s11225-009-9184-3

dimension $\delta \in \Delta$. Semantically, the structures are extended with one total relation for each dimension (over the same state space). Many applications can be envisioned for such a multi-modal variant of CTL, called MCTL. State transition systems are popular as formal models of *multi-agent systems* [25]. If we make the assumptions that agents can act whenever they want and never act at exactly the same time, we essentially have a structure where the transitions are labeled by agent names (and where there is at least one outgoing transition for each agent in each state), and a formula of the form $\mathsf{E}^a\phi$ means that if only agent a acts then she can act in such a way that ϕ is true. A related example is reasoning about interleaving computations of several processes with shared resources. MCTL may also find application as a query language over tree-like structures. For instance, take XPath, a language used to navigate through elements and attributes in an XML document [23]. Gottlob and Kock [15] use different versions of the tense modalities corresponding to the different directions in XPath in order to encode a fragment of XPath. We could, e.g., take $\mathsf{E}^{\downarrow}\bigcirc$ to mean "there is a next child" and $\mathsf{E}^{\rightarrow}\bigcirc$ to mean "there is a next sibling".

Another application closely related to multi-agent systems is that of *normative systems*. Typically, such a system defines a set of constraints on the behaviour of agents, corresponding to obligations and permissions. In Normative Temporal Logic (NTL), described in Section 5, obligations and permissions are, first, contextualised to a normative system η and, second, receive a temporal dimension. That is, expressions introduced in [1] of the forms $\mathsf{P}^\eta\alpha$ and $\mathsf{O}^\eta\alpha$ (where α starts with a tense modality, and η is a normative system), means that in the context of the normative system η, α is permitted or obligatory, respectively.

In the next section, CTL is briefly reviewed, before MCTL is formally defined in Section 3. The axiomatisation and completeness proof are found in Section 4. We first give an informal outline of the proof and a detailed example, before we describe the proof in detail in Section 4.2. Then, in Section 5 we present Normative Temporal Logic (NTL) as an instance of MCTL. In Section 6 we conclude. This paper includes material from the M4M proceedings [2] and the book chapter [3].

2. CTL

Given a set of primitive propositions Θ, the language $\mathcal{L}_{CTL}(\Theta)$ of CTL is defined by the following grammar.

$$\phi ::= \top \mid p \mid \neg\phi \mid \phi \vee \phi \mid \mathsf{E}\bigcirc\phi \mid \mathsf{E}(\phi\,\mathcal{U}\,\phi) \mid \mathsf{A}\bigcirc\phi \mid \mathsf{A}(\phi\,\mathcal{U}\,\phi)$$

where $p \in \Theta$. The usual derived propositional connectives are used, in addition to $\mathsf{E}\Diamond\phi$ ($\mathsf{A}\Diamond\phi$) for $\mathsf{E}(\top\,\mathcal{U}\,\phi)$ ($\mathsf{A}(\top\,\mathcal{U}\,\phi)$) and $\mathsf{E}\,\Box\phi$ ($\mathsf{A}\,\Box\phi$) for $\neg\mathsf{A}\Diamond\neg\phi$ ($\neg\mathsf{E}\Diamond\neg\phi$).

A CTL *model* over Θ is a tuple $M = (S, R, L)$ where S is a set of states, $R \subseteq S \times S$ is total[1] and $L(s) \subseteq \Theta$ for each $s \in S$. The class of all models over Θ is denoted $\mathcal{M}_{CTL}(\Theta)$. A model is *finite* if the set of states is finite. In general, given a set S and a total relation R over S, we will use $\pi(R, s)$ to denote the R-paths starting in s, i.e., the set of sequences $x_0x_1\cdots$ such that $x_0 = s$ and for each $i \geq 0$, $(x_i, x_{i+1}) \in R$. For $x \in \pi(R, s)$ and $k \geq 0$, $x[k]$ denotes the the kth element of x (x_k). A *pointed model* is a pair M, s where M is a model and s is a state in M. Satisfaction is defined as follows.

$$
\begin{array}{lll}
M, s \models_{CTL} \top & & \\
M, s \models_{CTL} p & \Leftrightarrow & p \in L(s) \quad (p \in \Theta) \\
M, s \models_{CTL} \neg\phi & \Leftrightarrow & M, s \not\models_{CTL} \phi \\
M, s \models_{CTL} \phi \vee \psi & \Leftrightarrow & M, s \models_{CTL} \phi \text{ or } M, s \models_{CTL} \psi \\
M, s \models_{CTL} \mathsf{E}\bigcirc\phi & \Leftrightarrow & \exists(x \in \pi(R, s))M, x[1] \models_{CTL} \phi \\
M, s \models_{CTL} \mathsf{A}\bigcirc\phi & \Leftrightarrow & \forall(x \in \pi(R, s))M, x[1] \models_{CTL} \phi \\
M, s \models_{CTL} \mathsf{E}(\phi\,\mathcal{U}\,\psi) & \Leftrightarrow & \exists(x \in \pi(R, s))\exists(j \geq 0)M, x[j] \models_{CTL} \psi \\
 & & \quad \text{and } \forall(0 \leq k < j)M, x[k] \models_{CTL} \phi \\
M, s \models_{CTL} \mathsf{A}(\phi\,\mathcal{U}\,\psi) & \Leftrightarrow & \forall(x \in \pi(R, s))\exists(j \geq 0)M, x[j] \models_{CTL} \psi \\
 & & \quad \text{and } \forall(0 \leq k < j)M, x[k] \models_{CTL} \phi
\end{array}
$$

Let $\mathcal{S}_{CTL}(\Theta)$ be the logical system over $\mathcal{L}_{CTL}(\Theta)$ defined in Figure 1. The following theorem gives completeness and decidability of CTL.

THEOREM 2.1 ([9]). *Any $\mathcal{S}_{CTL}(\Theta)$-consistent $\mathcal{L}_{CTL}(\Theta)$-formula is satisfiable in a finite $\mathcal{M}_{CTL}(\Theta)$ model.*

3. Multi-Modal CTL

We now define a multi-modal version of CTL. Let Δ be a finite set of indices and Θ a set of primitive propositions. The language $\mathcal{L}_{MCTL}(\Theta, \Delta)$ of MCTL is defined by the following grammar.

$$\phi ::= \top \mid p \mid \neg\phi \mid \phi \vee \phi \mid \mathsf{E}^{\delta}\bigcirc\phi \mid \mathsf{E}^{\delta}(\phi\,\mathcal{U}\,\phi) \mid \mathsf{A}^{\delta}\bigcirc\phi \mid \mathsf{A}^{\delta}(\phi\,\mathcal{U}\,\phi)$$

[1] For every $s \in S$ there is some $s' \in S$ such that Rss'.

(Ax1) All validities of propositional logic

(Ax4) $\mathsf{E}\bigcirc(\phi \vee \psi) \leftrightarrow (\mathsf{E}\bigcirc\phi \vee \mathsf{E}\bigcirc\psi)$

(Ax5) $\mathsf{A}\bigcirc\phi \leftrightarrow \neg\mathsf{E}\bigcirc\neg\phi$

(Ax6) $\mathsf{E}(\phi\,\mathcal{U}\,\psi) \leftrightarrow (\psi \vee (\phi \wedge \mathsf{E}\bigcirc\mathsf{E}(\phi\,\mathcal{U}\,\psi)))$

(Ax7) $\mathsf{A}(\phi\,\mathcal{U}\,\psi) \leftrightarrow (\psi \vee (\phi \wedge \mathsf{A}\bigcirc\mathsf{A}(\phi\,\mathcal{U}\,\psi)))$

(Ax8) $\mathsf{E}\bigcirc\top \wedge \mathsf{A}\bigcirc\top$

(Ax9) $\mathsf{A}\,\Box(\phi \rightarrow (\neg\psi \wedge \mathsf{E}\bigcirc\phi)) \rightarrow (\phi \rightarrow \neg\mathsf{A}(\gamma\,\mathcal{U}\,\psi))$

(Ax9b) $\mathsf{A}\,\Box(\phi \rightarrow (\neg\psi \wedge \mathsf{E}\bigcirc\phi)) \rightarrow (\phi \rightarrow \neg\mathsf{A}\Diamond\psi)$

(Ax10) $\mathsf{A}\,\Box(\phi \rightarrow (\neg\psi \wedge (\gamma \rightarrow \mathsf{A}\bigcirc\phi))) \rightarrow (\phi \rightarrow \neg\mathsf{E}(\gamma\,\mathcal{U}\,\psi))$

(Ax10b) $\mathsf{A}\,\Box(\phi \rightarrow (\neg\psi \wedge \mathsf{A}\bigcirc\phi)) \rightarrow (\phi \rightarrow \neg\mathsf{E}\Diamond\psi)$

(Ax11) $\mathsf{A}\,\Box(\phi \rightarrow \psi) \rightarrow (\mathsf{E}\bigcirc\phi \rightarrow \mathsf{E}\bigcirc\psi)$

(R1) If $\vdash \phi$ then $\vdash \mathsf{A}\,\Box\phi$ (generalization)

(R2) If $\vdash \phi$ and $\vdash \phi \rightarrow \psi$ then $\vdash \psi$ (modus ponens)

Figure 1. $\mathcal{S}_{CTL}(\Theta)$ [9]

where $\delta \in \Delta$ and $p \in \Theta$. The usual derived propositional connectives are used, in addition to $\mathsf{E}^\delta\Diamond\phi$ ($\mathsf{A}^\delta\Diamond\phi$) for $\mathsf{E}^\delta(\top\,\mathcal{U}\,\phi)$ ($\mathsf{A}^\delta(\top\,\mathcal{U}\,\phi)$) and $\mathsf{E}^\delta\,\Box\phi$ ($\mathsf{A}^\delta\,\Box\phi$) for $\neg\mathsf{A}^\delta\Diamond\neg\phi$ ($\neg\mathsf{E}^\delta\Diamond\neg\phi$).

We will use the following terminology: a *temporal atom* is a formula starting with a temporal operator; a *temporal δ-atom*, or sometimes just a δ-atom, is a formula starting with a temporal operator marked with δ.

Combinations of modal logics, e.g., of epistemic logic and temporal logic, have been studied to some extent both for particular logics and from a more abstract viewpoint [14]. Combinations of temporal logics into multi-dimensional temporal logics have been studied in the non-branching case [13], but we are not aware of existing results for similar combinations of branching-time logics such as CTL. Multi-modal CTL can be seen as a *fusion* of several "copies" of CTL. Studies of fusions and other combinations of modal logics have focussed on the transfer of meta-logical properties of the combined logics, such as soundness, completeness, decidability, etc. Many general transfer results exist for the fusion of *normal* modal logics [17, 11, 14]. However, CTL is not a normal modal logic[2], and these general results do not apply directly. Moreover, it is known that the common proof strategy of viewing the fusion as the union of *iterated modalisations* cannot always be used for non-normal modal logics [10]. The proof strategy we employ in this paper has similarities with the mentioned common strategy, but is not a direct application of it.

A MCTL *model* over Θ and Δ is a tuple $M = (S, \{R_\delta : \delta \in \Delta\}, L)$ where S is a set of states, $R_\delta \subseteq S \times S$ is total for each δ and $L(s) \subseteq \Theta$ for each $s \in S$. The class of all models over Θ and Δ is denoted $\mathcal{M}_{MCTL}(\Theta, \Delta)$.

The satisfaction relation between pointed $\mathcal{M}_{MCTL}(\Theta, \Delta)$ models and $\mathcal{L}_{MCTL}(\Theta, \Delta)$ formulae is defined exactly as for CTL, only that R_δ is used to interpret temporal operators marked with δ:

$$\begin{array}{lll}
M, s \models \mathsf{E}^\delta \bigcirc \phi & \Leftrightarrow & \exists (x \in \pi(R_\delta, s)) M, x[1] \models \phi \\
M, s \models \mathsf{A}^\delta \bigcirc \phi & \Leftrightarrow & \forall (x \in \pi(R_\delta, s)) M, x[1] \models \phi \\
M, s \models \mathsf{E}^\delta (\phi \mathcal{U} \psi) & \Leftrightarrow & \exists (x \in \pi(R_\delta, s)) \exists (j \geq 0) \\
& & M, x[j] \models \psi \text{ and } \forall (0 \leq k < j) M, x[k] \models \phi \\
M, s \models \mathsf{A}^\delta (\phi \mathcal{U} \psi) & \Leftrightarrow & \forall (x \in \pi(R_\delta, s)) \exists (j \geq 0) \\
& & M, x[j] \models \psi \text{ and } \forall (0 \leq k < j) M, x[k] \models \phi
\end{array}$$

4. Axiomatisation

Let $\mathcal{S}_{MCTL}(\Theta, \Delta)$ be the logical system over the language $\mathcal{L}_{MCTL}(\Theta, \Delta)$ obtained by taking one "copy" of the CTL axiomatisation for each dimension, as defined in Figure 2. We will show that $\mathcal{S}_{MCTL}(\Theta, \Delta)$ is sound and complete with respect to $\mathcal{M}_{MCTL}(\Theta, \Delta)$.

PROPOSITION 4.1. *$\mathcal{S}_{MCTL}(\Theta, \Delta)$ is sound wrt. $\mathcal{M}_{MCTL}(\Theta, \Delta)$.*

THEOREM 4.2. *Any $\mathcal{S}_{MCTL}(\Theta, \Delta)$-consistent $\mathcal{L}_{MCTL}(\Theta, \Delta)$-formula is satisfiable in a finite $\mathcal{M}_{MCTL}(\Theta, \Delta)$ model.*

The proof of Theorem 4.2 is presented in the following subsections. The following corollaries are immediate.

---

[2]To see that CTL is indeed not a normal modal logic, first observe that, e.g., $\mathsf{E}\Box$ neither distributes over conjunction nor disjunction and is thus neither a "box" nor a "diamond" of a normal modal logic. $\mathsf{E}\Box$ is derived, however, but we can make a similar argument for, e.g., the primary operator $\mathsf{A}\mathcal{U}$. Note that $\mathsf{A}\mathcal{U}$ is a dyadic operator; see [5, p. 195] for definitions of normality and the K axiom (and duals) for arbitrary similarity types. It is easy to see that the K axiom does not hold for the $\mathsf{A}\mathcal{U}$ operator (nor does it hold for the dual of that operator).

(Ax1) All validities of propositional logic

(Ax4) $\mathsf{E}^\delta \bigcirc (\phi \vee \psi) \leftrightarrow (\mathsf{E}^\delta \bigcirc \phi \vee \mathsf{E}^\delta \bigcirc \psi)$

(Ax5) $\mathsf{A}^\delta \bigcirc \phi \leftrightarrow \neg \mathsf{E}^\delta \bigcirc \neg \phi$

(Ax6) $\mathsf{E}^\delta(\phi \mathcal{U} \psi) \leftrightarrow (\psi \vee (\phi \wedge \mathsf{E}^\delta \bigcirc \mathsf{E}^\delta(\phi \mathcal{U} \psi)))$

(Ax7) $\mathsf{A}^\delta(\phi \mathcal{U} \psi) \leftrightarrow (\psi \vee (\phi \wedge \mathsf{A}^\delta \bigcirc \mathsf{A}^\delta(\phi \mathcal{U} \psi)))$

(Ax8) $\mathsf{E}^\delta \bigcirc \top \wedge \mathsf{A}^\delta \bigcirc \top$

(Ax9) $\mathsf{A}^\delta \Box(\phi \rightarrow (\neg\psi \wedge \mathsf{E}^\delta \bigcirc \phi)) \rightarrow (\phi \rightarrow \neg \mathsf{A}^\delta(\gamma \mathcal{U} \psi))$

(Ax9b) $\mathsf{A}^\delta \Box(\phi \rightarrow (\neg\psi \wedge \mathsf{E}^\delta \bigcirc \phi)) \rightarrow (\phi \rightarrow \neg \mathsf{A}^\delta \Diamond \psi)$

(Ax10) $\mathsf{A}^\delta \Box(\phi \rightarrow (\neg\psi \wedge (\gamma \rightarrow \mathsf{A}^\delta \bigcirc \phi))) \rightarrow (\phi \rightarrow \neg \mathsf{E}^\delta(\gamma \mathcal{U} \psi))$

(Ax10b) $\mathsf{A}^\delta \Box(\phi \rightarrow (\neg\psi \wedge \mathsf{A}^\delta \bigcirc \phi)) \rightarrow (\phi \rightarrow \neg \mathsf{E}^\delta \Diamond \psi)$

(Ax11) $\mathsf{A}^\delta \Box(\phi \rightarrow \psi) \rightarrow (\mathsf{E}^\delta \bigcirc \phi \rightarrow \mathsf{E}^\delta \bigcirc \psi)$

(R1) If $\vdash \phi$ then $\vdash \mathsf{A}^\delta \Box \phi$ (generalization)

(R2) If $\vdash \phi$ and $\vdash \phi \rightarrow \psi$ then $\vdash \psi$ (modus ponens)

Figure 2. $\mathcal{S}_{MCTL}(\Theta, \Delta)$. δ ranges over Δ.

COROLLARY 4.3. *$\mathcal{S}_{MCTL}(\Theta, \Delta)$ is complete wrt. $\mathcal{M}_{MCTL}(\Theta, \Delta)$.*

COROLLARY 4.4. *The satisfiability problem for* MCTL *is decidable.*

In fact, we can sharpen this result: we will show that, as a corollary of the construction used in the proof of Theorem 4.2, the satisfiability problem is in fact decidable in exponential time (and is thus EXPTIME-complete — no harder than the corresponding problem for CTL).

4.1. Outline of Completeness Proof

Let ϕ_0 be a consistent formula. Rather than extending the tableau-based method for proving the completeness of CTL in [8], we use a construction which employs the CTL completeness result (Theorem 2.1) directly, viewing a formula as a CTL formula for one dimension $\delta \in \Delta$ at a time by reading A^δ and E^δ as CTL path quantifiers A and E, respectively, and treating formulae starting with a δ'-operator ($\delta' \neq \delta$) as atomic formulae. By completeness of CTL, we get a CTL model for the formula (if it is consistent), where the states are labelled with atoms such as $\mathsf{A}^{\delta'}\psi$ or $\mathsf{E}^{\delta'}\psi$ (for $\delta' \neq \delta$). Then, for each δ' and each state, we expand the state by taking the conjunction of δ'-formulae the state is labelled with, construct a (uni-modal) CTL model of that formula, and "glue" the root of the model together with the state. Repeat for all dimensions and all states.

In order to keep the formulae each state is labelled with finite, we consider only subformulae of ϕ_0; by a δ-atom we here mean a subformula of ϕ_0 starting with either E^δ or A^δ. Let $At^{-\delta}$ denote the union of all sets of δ'-atoms for each $\delta' \neq \delta$. Furthermore, we assume that ϕ_0 is such that every occurrence of $\mathsf{E}^\delta(\alpha_1 \mathcal{U} \alpha_2)$ $(\mathsf{A}^\delta(\alpha_1 \mathcal{U} \alpha_2))$ is immediately preceeded by $\mathsf{E}^\delta \bigcirc$ $(\mathsf{A}^\delta \bigcirc)$ — we call this *XU form.* Any formula can be rewritten to XU form by recursive use of the axioms (Ax6) and (Ax7). We start with a model with a single state labelled with the literals in a consistent disjunct of ϕ_0 written in disjunctive normal form. We continue by expanding states labelled with formulae, one dimension δ at a time. In general, let $at(\delta, s)$ be the union of the set of δ-atoms s is labelled with and the set of negated δ-atoms of XU form s is *not* labelled with. We can now view $\bigwedge at(\delta, s)$ as a CTL formula over a language with primitive propositions $\Phi \cup At^{-\delta}$. The following can be shown: any MCTL consistent formula is satisfied by a state s' in some finite CTL model M' viewing $\Phi \cup At^{-\delta}$ as primitive propositions, such that for any $\delta' \neq \delta$ and any state t of M', $\bigwedge at(\delta', t)$ is MCTL-consistent, and s' does not have any ingoing transitions. This ensures that we can "glue" the pointed model M', s' to the state s while labelling the transitions in the model with the dimension δ we expanded — M', s' satisfies the formulae needed to be true there. The fact that s' does not have any ingoing transitions ensures that we can append M', s' to s without changing the truth of δ-atoms at s'. The fact that ϕ_0 is of XU form ensures that all labelled formulae are of XU form, which again ensures that we don't add new labels to a state when we expand it (because all the formulae we expand start with a next-modality). The fact that $\bigwedge at(\delta', t)$ is consistent for states t in the expanded model, ensures that we can repeat the process. Only a finite number of repetitions are needed, depending on the number of nested operators of different dimensions in ϕ_0, after which we can remove the non-Φ labels without affecting the truth of ϕ_0 and obtain a proper model.

4.1.1. Example

Take $\Delta = \{a, b\}$ and $\Theta = \{p, q, r\}$. We illustrate the method for finding a satisfying MCTL model for the formula

$$\phi_0 = \mathsf{E}^a \bigcirc (p \wedge \mathsf{E}^b \bigcirc (q \wedge \mathsf{E}^a \bigcirc r) \wedge \mathsf{E}^a(r \mathcal{U} \neg p) \wedge \mathsf{A}^a \bigcirc p) \wedge \mathsf{A}^a \bigcirc q \wedge \mathsf{E}^b \bigcirc p$$

We define the model in steps. Some of the information given here for each step refers to the proof in the following section.

The initial model M_0 consists simply of a single state $\hat{s}$ labelled with the temporal atoms required to be true. In this model *every* temporal atom is viewed as a primitive proposition.

$\mathbf{M^0}$: $(U^0 = \{\hat{s}\},\ T^0 = \emptyset,\ \tau^0(\hat{s}) = \epsilon)$

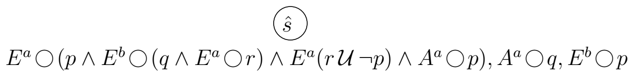

In general, the model M_{j+1} is constructed from M_j by expanding each node in U^j by constructing one CTL model for the temporal atoms in that node of each dimension, and then attaching these CTL models to the node we expand.

Expanding $\hat{s}$ along dimension a, we treat the temporal atoms of a dimension different from a as primitive propositions, and E^a and A^a as the CTL path quantifiers E and A, respectively. From completeness of CTL we know that there is a model for the formulae $\hat{s}$ is labelled with. There are, of course, *many* CTL models, but we choose a model where the labels (temporal atoms of dimensions different from a) are MCTL-*consistent* — which ensures that we can repeat the process and expand the new nodes again by choosing a CTL model — and where there are no ingoing transitions to the root — ensuring that we can glue models of different dimensions together. (The existence of models with these properties is formally ensured by Proposition 4.8 below). We get, e.g., the following (uni-modal) CTL-model (right), satisfying the set of CTL formulae $\{\mathsf{E}(p \wedge t \wedge \mathsf{E}(r\, \mathcal{U} \neg p) \wedge A \bigcirc p), \mathsf{A} \bigcirc q\}$, where t is an atom representing $\mathsf{E}^b \bigcirc (q \wedge \mathsf{E}^a \bigcirc r)$. This is a proper CTL-model, with a single, total, relation. Expanding $\hat{s}$ along dimension b we get the (uni-modal) CTL-model on the left:

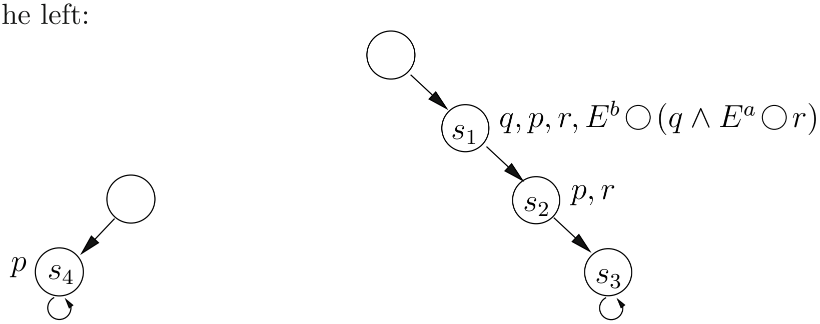

There was only one state in U^0, and two dimensions, so we are done. Gluing the two CTL models together with the state $\hat{s}$ we expanded, we get M_1:

M_1: $(U^1 = \{s_1, s_2, s_3, s_4\},\ T^1 = \{\hat{s}\},\ \tau^1(s_1) = \tau^1(s_2) = \tau^1(s_3) = a; \tau^1(s_4) = b)$

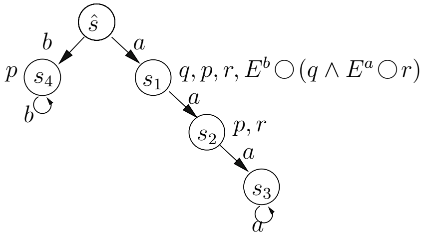

U^1 is the set of nodes added in the previous round, which will be expanded now. It might seem that s_4 does not need to be expanded because it is not labelled by any temporal formulae, but it must be expanded along the a-dimension in a trivial way: a self loop must be added to make sure that the a-relation is total. Similarly for s_2 and s_3 wrt. b. The result is M_2:

M_2: $(U^2 = \{s_5\}, T^2 = \{\hat{s}, s_1, s_2, s_3, s_4\}, \tau^2(s_5) = b)$

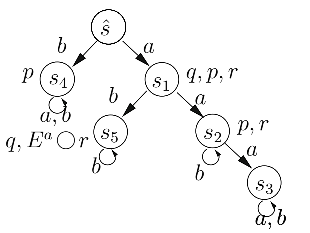

M_3 is as follows:

M_3: $(U^3 = \{s_6\}, T^3 = \{\hat{s}, s_1, s_2, s_3, s_4, s_5\}, \tau^3(s_6) = a)$

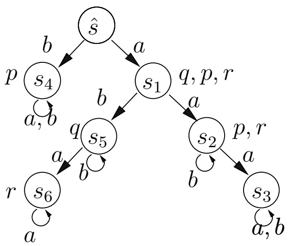

Finally, M_4 trivially expands s_6 by gluing on a model for each of the formulae in each of the dimensions different from $\tau(s_6)$. There are no such formulae, so the models are trivial (satisfying tautologies) but total:

M_4: $(U^4 = \emptyset,\ T^4 = \{\hat{s}, s_1, s_2, s_3, s_4, s_5, s_6\})$

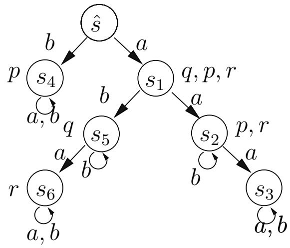

There are no more states to expand, and the construction is finished.

4.2. Completeness Proof

We now formally prove Theorem 4.2.

Let ϕ_0 be a $\mathcal{S}_{MCTL}(\Theta, \Delta)$ consistent $\mathcal{L}_{MCTL}(\Theta, \Delta)$ formula. We will show that ϕ_0 is satisfied by a finite model in $\mathcal{M}_{MCTL}(\Theta, \Delta)$. We repeat the definition of XU form:

Definition 4.5 (XU form). A formula $\phi \in \mathcal{L}_{MCTL}(\Theta, \Delta)$ is of *XU form* if every occurrence of a subformula of the form $\mathsf{E}^\delta(\psi_1 \mathcal{U} \psi_2)$ $(\mathsf{A}^\delta(\psi_1 \mathcal{U} \psi_2))$ in ϕ is immediately preceeded by an $\mathsf{E}^\delta \bigcirc$ $(\mathsf{A}^\delta \bigcirc)$ operator.

Lemma 4.6. *Any $\mathcal{L}_{MCTL}(\Theta, \Delta)$ formula ϕ is equivalent to a $\mathcal{L}_{MCTL}(\Theta, \Delta)$ formula of XU form.*

Proof. Rewrite the formula using axioms (Ax6) and (Ax7) (which are easily seen to be valid) recursively, until the formula is of the form. ■

Thus, we will henceforth assume that ϕ_0 is of XU form. Let $Subf(\phi)$ be the set of all subformulae of a formula ϕ. We can view the language $\mathcal{L}_{MCTL}(\Theta, \Delta)$ as a CTL language, by fixing some δ and reading $\mathsf{E}^\delta X$ as $\mathsf{E}X$, $\mathsf{A}^\delta X$ as $\mathsf{A}X$, and so on, and treating the other temporal atoms, such as $\mathsf{E}^{\delta'} X\phi$, $\delta' \neq \delta$, as primitive propositions (in addition to Θ). For technical reasons, we only consider temporal atoms occurring in $Subf(\phi_0)$. Let:

$$At^\delta = \{\mathsf{E}^\delta \bigcirc \phi, \mathsf{E}^\delta(\phi \mathcal{U} \psi), \mathsf{A}^\delta \bigcirc \phi, \mathsf{A}^\delta(\phi \mathcal{U} \psi) : \phi, \psi \in Subf(\phi_0)\}$$

– in particular, At^δ includes the set of temporal atoms of type δ occurring in ϕ_0 — let

$$At = \bigcup_{\delta \in \Delta} At^\delta$$

– hence, At includes all temporal atoms in ϕ_0 — and let

$$At^{-\delta} = \bigcup_{\delta' \neq \delta} At^{\delta'}$$

– thus, $At^{-\delta}$ includes the temporal atoms occurring in ϕ_0 which are not of type δ. We can now view any formula in $\mathcal{L}_{MCTL}(\Theta, \Delta) \cap Subf(\phi_0)$ as a $\mathcal{L}_{CTL}(At^{-\delta} \cup \Theta)$ formula by reading any $\mathsf{E}^{\delta}, \mathsf{A}^{\delta}$ which is not in the scope of any $\mathsf{E}^{\delta'}, \mathsf{A}^{\delta'}$ ($\delta' \neq \delta$) as E, A, and treating temporal formulae such as $\mathsf{E}^{\delta'} X\phi$ where $\delta' \neq \delta$ as primitive propositions. When Θ and Δ are understood, we will use $\mathcal{L}_{CTL}(\delta)$ as shorthand for the CTL language $\mathcal{L}_{CTL}(At^{-\delta} \cup \Theta)$ and $\mathcal{M}_{CTL}(\delta)$ as a shorthand for the associated CTL model class $\mathcal{M}_{CTL}(At^{-\delta} \cup \Theta)$. A model $M \in \mathcal{M}_{CTL}(\delta)$ has a transition relation for interpreting temporal δ-atoms, and the labelling function interprets the other temporal atoms occurring in ϕ_0 in addition to primitive propositions Θ in the states. Similarly, we use $\mathcal{S}_{CTL}(\delta)$ to denote the CTL axiom system $\mathcal{S}_{CTL}(At^{-\delta} \cup \Theta)$ over the language $\mathcal{L}_{CTL}(\delta)$. Thus, we will henceforth sometimes view a MCTL formula ϕ also as a $\mathcal{L}_{CTL}(\delta)$ formula for some given δ, and write, e.g., $M, s \models_{CTL} \phi$ when $M \in \mathcal{M}_{CTL}(\delta)$ with the meaning defined by reading E^{δ} as E, etc., as explained above. Similarly, we sometimes implicitly view a $\mathcal{L}_{CTL}(\delta)$ formula as a MCTL formula (i.e., the MCTL formula obtained by replacing every E with E^{δ} and every A with A^{δ}).

LEMMA 4.7. *For any δ and $\phi \in \mathcal{L}_{CTL}(\delta)$, $\vdash_{\mathcal{S}_{CTL}(\delta)} \phi$ implies that $\vdash_{\mathcal{S}_{MCTL}} \phi$.*

PROOF. Straightforward induction on the length of the proof. ■

When t is a state of a model $M \in \mathcal{M}_{CTL}(\delta)$ and $\delta' \neq \delta$, let

$$\begin{aligned} at(\delta', t, M) = \quad & \{\psi : \psi \in At^{\delta'}, \psi \text{ is of XU form}, \psi \in L(t)\} \cup \\ & \{\neg\psi : \psi \in At^{\delta'}, \psi \text{ is of XU form}, \psi \notin L(t)\} \end{aligned}$$

PROPOSITION 4.8. *Let $\delta \in \Delta$ and $\phi \in \mathcal{L}_{CTL}(\delta)$. If ϕ is $\mathcal{S}_{MCTL}$-consistent, then there is a model $M' \in \mathcal{M}_{CTL}(\delta)$ with a state s' such that*

1. $M', s' \models_{CTL} \phi$
2. *For all states t reachable from s' in M' and for all $\delta' \neq \delta$, $\bigwedge at(\delta', t, M')$ is $\mathcal{S}_{MCTL}$-consistent*
3. *There is no state t in M' such that $(t, s') \in R'$*
4. *M' is finite*

Proof. Let $XU^{\delta'}$ be the set of all formulae in $At^{\delta'}$ of XU form, and let $XU^{\delta'+}$ be $XU^{\delta'}$ closed under single negation, i.e., $XU^{\delta'+} = \{\alpha, \neg\alpha : \alpha \in At^{\delta'}, \alpha \text{ of XU form}\}$. Let $Y^{\delta'}$ be the set of all $XU^{\delta'+}$-maximal $\mathcal{S}_{MCTL}$-inconsistent subsets of $XU^{\delta'+}$, i.e., all sets $y \subseteq XU^{\delta'+}$ such that either $\alpha \in y$ or $\neg\alpha \in y$ for any $\alpha \in XU^{\delta'}$ and $\vdash_{\mathcal{S}_{MCTL}} \bigwedge y \rightarrow \bot$. $Y^{\delta'}$ is finite because $XU^{\delta'+}$ is finite. Let

$$f(\delta') = \bigwedge y_1 \vee \cdots \vee \bigwedge y_k \qquad \text{where } Y^{\delta'} = \{y_1, \ldots, y_k\}$$

We show that

$$\gamma = \phi \wedge \mathsf{A}^\delta \Box \bigwedge_{\delta' \neq \delta} \neg f(\delta')$$

is $\mathcal{S}_{MCTL}$-consistent. Assume the opposite: $\vdash_{\mathcal{S}_{MCTL}} \gamma \rightarrow \bot$. It follows that $\vdash_{\mathcal{S}_{MCTL}} \mathsf{A}^\delta \Box \bigwedge_{\delta' \neq \delta} \neg f(\delta') \rightarrow \neg\phi$. However, for any $\delta' \neq \delta$ and $y \in Y^{\delta'}$ we have that $\vdash_{\mathcal{S}_{MCTL}} \neg \bigwedge y$, and thus that $\vdash_{\mathcal{S}_{MCTL}} \neg f(\delta')$ for any δ'. It follows that $\vdash_{\mathcal{S}_{MCTL}} \bigwedge_{\delta' \neq \delta} \neg f(\delta')$. By (Gen), we have that $\vdash_{\mathcal{S}_{MCTL}} \mathsf{A}^\delta \Box \bigwedge_{\delta' \neq \delta} \neg f(\delta')$. But then we also have that $\vdash_{\mathcal{S}_{MCTL}} \neg\phi$, which contradicts the fact that ϕ is $\mathcal{S}_{MCTL}$-consistent. Thus, γ is $\mathcal{S}_{MCTL}$-consistent.

Clearly, γ is $\mathcal{S}_{CTL}(\delta)$-consistent — otherwise it would not have been $\mathcal{S}_{MCTL}$-consistent by Lemma 4.7. By completeness of $\mathcal{S}_{CTL}(\delta)$ (Theorem 2.1), there is a finite model $M = (S, R, L) \in \mathcal{M}_{CTL}(\delta)$ such that $M, s \models_{CTL} \gamma$ for some s. Let t be reachable from s in M. Assume that $\bigwedge at(\delta', t, M)$ is not $\mathcal{S}_{MCTL}$-consistent for some $\delta' \neq \delta$. Then $at(\delta', t, M) = y_j$ for some j, so $M, t \models_{CTL} f(\delta')$. It follows that $M, s \models_{CTL} \mathsf{E}^\delta \Diamond f(\delta')$, but this contradicts the fact that $M, s \models_{CTL} \gamma$. Thus, $\bigwedge at(\delta', t, M)$ is $\mathcal{S}_{MCTL}$-consistent. Also, $M, s \models_{CTL} \phi$.

To get a satisfying state with no ingoing transitions, let $M' = (S', R', L')$ where $S' = S \cup \{s'\}$ for some new state s'; $R' = R \cup \{(s', t) : (s, t) \in R\}$; $L'(s') = L(s)$ and $L'(t) = L(t)$ for $t \neq s'$. It is easy to see that $M, s \models_{CTL} \psi$ iff $M', s' \models_{CTL} \psi$ for all ψ. In particular $M', s' \models_{CTL} \phi$. ■

Definition 4.9 (General Models). A *general model* over Θ and Δ is a tuple $M = (S, T, U, \tau, \{R_\delta : \delta \in \Delta\}, L, K)$ where T and U partition S, $\tau(u) \in \Delta \cup \{\epsilon\}$ for each $u \in U$, $K(u) \subseteq \bigcup_{\delta' \neq \tau(u)} At^{\delta'}$ for each $u \in U$, and the other elements are as in a model. A general model is finite if S finite.

Satisfaction of a formula $\phi \subseteq Subf(\phi_0)$ in a pointed general model is defined as follows. Let $s \in S$.

$$\begin{array}{lcl}
M, s \models \top & & \\
M, s \models p & \Leftrightarrow & p \in L(s) \quad (p \in \Theta) \\
M, s \models \neg\phi & \Leftrightarrow & M, s \not\models \phi \\
M, s \models \phi \vee \psi & \Leftrightarrow & M, s \models \phi \text{ or } M, s \models \psi \\
M, s \models \mathsf{A}^\delta X\phi & \Leftrightarrow & \begin{cases} \forall (x \in \pi(R_\delta, s)) M, x[1] \models \phi & s \in T \\ \forall (x \in \pi(R_\delta, s)) M, x[1] \models \phi & s \in U \text{ and } \delta = \tau(s) \\ \mathsf{A}^\delta X\phi \in K(s) & s \in U \text{ and } \delta \neq \tau(s) \end{cases}
\end{array}$$

and similarly for the other temporal atoms.

We now define a sequence $M_0, M_1, \dots$ of finite general models $M_j = (S^j, T^j, U^j, \tau^j, \{R^j_\delta : \delta \in \Delta\}, L^j, K^j)$ such that $\hat{s} \in S^j$ for all j for some state $\hat{s}$, having the three following properties for any j:

(i) $M_j, \hat{s} \models \phi_0$

(ii) For every $t \in U^j$ and $\delta \neq \tau^j(t)$, $\bigwedge at(\delta, t, M_j)$ is $\mathcal{S}_{MCTL}$-consistent

(iii) For every $t \in U^j$, each $\alpha \in K^j(t)$ is of XU form

where

$$\begin{aligned} at(\delta, s, M_j) = \quad & \{\psi : \psi \in At^\delta, \psi \text{ is of XU form}, \psi \in K^j(s)\} \cup \\ & \{\neg\psi : \psi \in At^\delta, \psi \text{ is of XU form}, \psi \notin K^j(s)\} \end{aligned}$$

It might be instructive to refer to the example in the previous section as an illustration of the construction.

M_0 has a single state $\hat{s}$, such that $\hat{s} \in U^0$ and $\tau(\hat{s}) = \epsilon$. If we view *every* temporal atom in ϕ_0 which is not in the scope of another temporal operator as a primitive proposition, ϕ_0 is a purely propositional formula. Because $\mathcal{S}_{MCTL}$ contains propositional logic and ϕ_0 is of XU form, ϕ_0 is equivalent to a formula on disjunctive normal form $(A^1_1 \wedge \cdots \wedge A^1_m) \vee \cdots \vee (A^k_1 \wedge \cdots A^k_m)$, where for each $1 \leq j \leq k$ and $1 \leq i \leq m$, either $A^j_i = B_i$ or $A^j_i = \neg B_i$, where $\{B_1, \dots, B_m\} = \Theta \cup \{\alpha \in At : \alpha \text{ of XU form}\}$. Since ϕ_0 is $\mathcal{S}_{MCTL}$-consistent, some $\xi = (A^j_1 \wedge \cdots \wedge A^j_m)$ is $\mathcal{S}_{MCTL}$-consistent. Let X be the set of positive atoms A^j_i in ξ, and let $Y = \{B_1, \dots, B_m\} \setminus X$ be the negative atoms. I.e, $\xi = \bigwedge(X \cup \{\neg y : y \in Y\})$ is $\mathcal{S}_{MCTL}$-consistent. Set $L^0(\hat{s}) = X \cap \Theta$ and $K^0(\hat{s}) = X \setminus L^0(\hat{s})$. (i) clearly holds, because M^0 interprets ϕ_0 simply as a propositional formula using the valuations $L^0(\hat{s})$ and $K^0(\hat{s})$ and thus we see immediately that $M_0, \hat{s} \models \xi$. (ii) holds, because $at(\delta, \hat{s}, M_0) \subseteq X \cup \{\neg y : y \in Y\}$. (iii) holds immediately, because every atom in X is of XU form.

M_{j+1} is obtained from M_j as follows. Informally, the idea is to take, for every δ and every state s in U^j, the set $at(\delta, s, M_j)$, and replace it with a model $M \in \mathcal{M}_{CTL}(\delta)$ for $at(\delta, s, M_j)$ rooted in s. Formally, we define M_{j+1} as follows. For every $u \in U^j$ and every $\delta \neq \tau^j(u)$, we have that $\bigwedge at(\delta, u, M_j)$ is $\mathcal{S}_{MCTL}$-consistent by (ii), so take $\phi = \bigwedge at(\delta, u, M_j)$ and let $M' = (S', R', L') \in \mathcal{M}_{CTL}(\delta)$ and s' be as in Proposition 4.8. W.l.o.g. we assume that every state in M' is reachable from s'. Let $M_{u,\delta} = (S_{u,\delta}, R_{u,\delta}, L_{u,\delta}) \in \mathcal{M}_{CTL}(\delta)$ be equal to M' except that $L_{u,\delta}(s') = \emptyset$ and $L_{u,\delta}(t) = \{\alpha \in L'(t) : \alpha \text{ of XU form}\}$ for any $t \in S' \setminus \{s'\}$, and let $t_{u,\delta} = s'$. We have that each $\alpha \in at(\delta, u, M_j)$ starts with a (possibly negated) $\mathsf{E}^{\delta'} \bigcirc$ or $\mathsf{A}^{\delta'} \bigcirc$ operator (for some δ'). Together with the fact that s' does not have any ingoing transitions, this implies that changing $L'(s')$ does not affect the truth of $\bigwedge at(\delta, u, M_j)$ in s'. Furthermore, removing atoms not of XU form from $L'(t)$ does not affect the truth of $\bigwedge at(\delta, u, M_j)$ in s' either, because all atoms in $at(\delta, u, M_j)$ are of XU form. Thus, all the four points in Proposition 4.8 still hold for $M_{u,\delta}$ and $t_{u,\delta}$; in particular we have that $M_{u,\delta}, t_{u,\delta} \models_{CTL} \bigwedge at(\delta, u, M_j)$.

Let

$$
\begin{aligned}
&S^{j+1} = S^j \cup \textstyle\bigcup_{u \in U^j, \delta \neq \tau^j(u)} (S_{u,\delta} \setminus \{t_{u,\delta}\})\\
&T^{j+1} = T^j \cup U^j\\
&U^{j+1} = \textstyle\bigcup_{u \in U^j, \delta \neq \tau^j(u)} (S_{u,\delta} \setminus \{t_{u,\delta}\})\\
&\tau^{j+1}(v) = \delta \text{ iff } v \in S_{u,\delta} \text{ (for some } u)\\
&R_\delta^{j+1} = R_\delta^j \cup \textstyle\bigcup_{u \in U^j, \delta \neq \tau^j(u)} \{(x', y) : (x, y) \in R_{u,\delta}\}\\
&\text{where } x' = u \text{ if } x = t_{u,\delta} \text{ and } x' = x \text{ otherwise}\\
&L^{j+1}(t) = \begin{cases} L^j(t) & t \in S^j \\ L_{u,\delta}(t) \cap \Theta & t \in S_{u,\delta} \end{cases}\\
&K^{j+1}(t) = L_{u,\delta}(t) \setminus \Theta \text{ when } t \in S_{u,\delta}
\end{aligned}
$$

Since S^j is finite and each $S_{u,\delta}$ is finite (guaranteed by Proposition 4.8) and both U^j and Δ are finite, S^{j+1} is finite.

We argue that (ii) holds for M_{j+1}. Let $t \in U^{j+1}$ and $\delta' \neq \tau^{j+1}(t)$. $t \in S_{u,\delta}$ for some $u \in U^j$ and some $\delta \neq \tau^j(u)$, which implies that $\tau^{j+1}(t) = \delta$ and thus that $\delta \neq \delta'$. We have that $\bigwedge at(\delta', t, M_{j+1}) = \bigwedge at(\delta', t, M_{u,\delta})$ is $\mathcal{S}_{MCTL}$-consistent by Proposition 4.8.

We argue that (iii) holds for M_{j+1}. Let $t \in U^{j+1}$. We have that $K^{j+1}(t) = L_{u,\delta}(t) \setminus \Theta$, for some u, δ. (iii) holds immediately, because every formula in $L_{u,\delta}$ is of XU form by construction.

We now argue that also (i) holds for M_{j+1}. First we show that for any $v \in U^j$ and any $\beta \in Subf(\phi_0)$ of XU form

$$M_j, v \models \beta \Leftrightarrow M_{j+1}, v \models \beta \tag{I}$$

by induction on the structure of β.

- $\beta = p$: $L^j(v) = L^{j+1}(v)$.
- $\beta \in \{\mathsf{E}^{\delta'} \bigcirc \gamma, \mathsf{E}^{\delta'} \bigcirc \mathsf{E}^{\delta'}(\gamma_1 \mathcal{U} \gamma_2), \mathsf{A}^{\delta'} \bigcirc \gamma, \mathsf{A}^{\delta'} \bigcirc \mathsf{E}^{\delta'}(\gamma_1 \mathcal{U} \gamma_2) : \gamma, \gamma_1, \gamma_2$ of XU form$\}$: First assume that $\delta' \neq \tau^j(v)$. $M_j, v \models \beta$ iff $\beta \in K^j(v)$ iff $\beta \in at(\delta', v, M_j)$ iff $M_{v,\delta'}, t_{v,\delta'} \models \beta$ (the fact that $at(\delta', v, M_j)$ is closed under single negation gives us both directions). From the construction of $R^{\delta'}_{j+1}$, the only δ'-transitions from v are transitions from $t_{v,\delta'}$ in $R_{v,\delta}$: we have that
$$(t_{v,\delta'}, t) \in R_{v,\delta'} \Leftrightarrow (v, t) \in R^{\delta'}_{j+1}$$
for any t. Furthermore, we have that
$$M_{v,\delta'}, t \models \alpha \Leftrightarrow M_{j+1}, t \models \alpha$$
for any t such that $(v, t) \in R^{\delta'}_{j+1}$ and for any α: the submodel of M_{j+1} generated by t is equivalent to the submodel of $M_{v,\delta'}$ generated by t — this holds because $t_{v,\delta}$ does not have any ingoing transitions and thus v does not have any ingoing δ'-transitions — and these two submodels interpret formulae in exactly the same way. It follows that $M_{v,\delta'}, t_{v,\delta'} \models \beta$ iff $M_{j+1}, v \models \beta$.
Second, assume that $\delta' = \tau(v)$. Observe that $\pi(R^j_{\delta'}, v) = \pi(R^{j+1}_{\delta'}, v)$, and for any state w along a path we have that $w \in U^j$ by construction since $v \in U^j$, so $M_j, w \models \alpha$ iff $M_{j+1}, w \models \alpha$ for $\alpha \in \{\gamma, \gamma_1, \gamma_2\}$ by the induction hypothesis. It follows that $M_j, v \models \beta$ iff $M_{j+1}, v \models \beta$.
- Propositional connectives: Straightforward.

We now argue that for any $v \in S_j$ and any $\psi \in Subf(\phi_0)$ of XU form,

$$M_j, v \models \psi \Leftrightarrow M_{j+1}, v \models \psi$$

That (i) holds for M_{j+1} follows immediately. We argue by structural induction on ψ. For $\psi \in \Theta$, we have that $M_j, v \models \psi \Leftrightarrow M_{j+1}, v \models \psi$ because $L^j(v) = L^{j+1}(v)$. Assume that ψ is a temporal atom in the set $\{\mathsf{E}^{\delta} \bigcirc \gamma, \mathsf{E}^{\delta} \bigcirc \mathsf{E}^{\delta}(\gamma_1 \mathcal{U} \gamma_2), \mathsf{A}^{\delta} \bigcirc \gamma, \mathsf{A}^{\delta} \bigcirc \mathsf{A}^{\delta}(\gamma_1 \mathcal{U} \gamma_2) : \gamma, \gamma_1, \gamma_2$ of XU form$\}$, and consider first the case that $v \in T_j$. For any δ, the δ-paths in M_j starting in v are exactly the same as the δ-paths in M_{j+1} starting in v.

By the induction hypothesis, we have that $M_j, w \models \gamma$ iff $M_{j+1}, w \models \gamma$ for any w along any of these paths and any $\gamma \in Subf(\psi)$ of XU form, which shows that $M_j, v \models \psi$ iff $M_{j+1}, v \models \psi$. Consider, second, that $v \in U_j$, in which case we immediately have the required result by (I). The cases for the propositional connectives are straightforward. This concludes the argument that (i) holds for M_{j+1}.

Let the degree of a formula α, denoted $deg(\alpha)$, be the maximum number of nested temporal operators *of alternating type* in the formula. For example, $deg(p) = 0$, $deg(\mathsf{E}^\delta \bigcirc p) = 1$, $deg(\mathsf{E}^\delta \bigcirc \mathsf{A}^{\delta'}(p \mathcal{U} q)) = 2$ (two dimensions of alternating type), $deg(\mathsf{E}^\delta \bigcirc \mathsf{A}^\delta(p \mathcal{U} q)) = 1$ (two temporal operators but not of alternating type), $deg((\mathsf{E}^\delta \bigcirc p) \wedge (\mathsf{E}^{\delta'} \bigcirc q)) = 1$ (two dimensions but not nested), $deg(\mathsf{E}^\delta \bigcirc \mathsf{E}^{\delta'}(p \mathcal{U} \mathsf{A}^{\delta'} \bigcirc \mathsf{E}^\delta \bigcirc q)) = 4$, etc. Formally we can define $deg(\alpha)$ as follows. Let $dim(\alpha)$ be the set of dimensions of all occurrences of temporal operators in the formula α which are not in the scope of any other temporal operator: $dim(p) = \emptyset$; $dim(\neg\phi) = dim(\phi)$; $dim(\phi_1 \vee \phi_2) = dim(\phi_1) \cup dim(\phi_2)$; $dim(\mathsf{E}^\delta \bigcirc \phi) = dim(\mathsf{A}^\delta \bigcirc \phi) = dim(\mathsf{E}^\delta(\phi_1 \mathcal{U} \phi_2)) = dim(\mathsf{A}^\delta(\phi_2 \mathcal{U} \phi_2)) = \{\delta\}$. Finally, $deg(p) = 0$; $deg(\neg\phi) = deg(\phi)$; $deg(\phi_1 \vee \phi_2) = max(deg(\phi_1), deg(\phi_2))$;

$$deg(\mathsf{E}^\delta \bigcirc \phi) = deg(\mathsf{A}^\delta \bigcirc \phi) = \begin{cases} deg(\phi) & dim(\phi) \setminus \{\delta\} = \emptyset \\ deg(\phi) + 1 & \text{otherwise} \end{cases}$$

(increase the degree whenever there is a dimension different from δ in ϕ);

$$deg(\mathsf{E}^\delta(\phi_1 \mathcal{U} \phi_2)) = deg(\mathsf{A}^\delta(\phi_2 \mathcal{U} \phi_2)) = \begin{cases} deg(\phi) & dim(\phi) \setminus \{\delta\} = \emptyset \\ deg(\phi) + 1 & \text{otherwise} \end{cases}$$

where $\phi = \phi_1$ if $deg(\phi_1) > deg(\phi_2)$ and $\phi = \phi_2$ otherwise.

A general model $M = (S, T, U, \tau, \{R_\delta : \delta \in \Delta\}, L, K)$ is a generalisation of a proper model $M' = (S, \{R_\delta : \delta \in \Delta\}, L)$. We say that the satisfaction relationship between a formula ϕ and a pointed general model (M, s), i.e., the question of whether $M, s \models \phi$ or not, is *classical* if the definition (as given recursively above) does not involve any state from U (and thus not τ or K either). If the satisfaction relationship is classical, then satisfaction only depends on the underlying (proper) model.

LEMMA 4.10. *For any $\alpha \in Subf(\phi_0)$, the satisfaction relationship between α and (M_{m+j+1}, v) is classical when $deg(\alpha) = j$ and $v \in U^m$.*

PROOF. Directly from the definition of satisfaction, we have that when $v \in U^m$ then for every $k > m$ if $x \in \pi(R^k_\delta, v)$ for some δ then

1. if $\tau^m(v) = \delta$, then for any i, $x[i] \in U^m$

2. if $\tau^m(v) \neq \delta$, then for any i, $x[i] \in U^m \cup U^{m+1}$

This means that when we evaluate a formula in a state $v \in U^m$, only a "switch" in dimension can involve states from U^{m+1}. For example, in the evaluation the formula $\mathsf{E}^\delta(\mathsf{A}^\delta \bigcirc p \mathcal{U} \mathsf{E}^\delta \bigcirc q)$ of degree 1 in a state $v \in U^m$ in a model M_k where $k > m$, only states $u \in U^m \cup U^{m+1}$ are involved. In the evaluation of the degree 2 formula $\mathsf{E}^\delta(\mathsf{A}^\delta \bigcirc p \mathcal{U} \mathsf{E}^{\delta'} \bigcirc q)$, only states in $U^m \cup U^{m+1} \cup U^{m+2}$ are involved. If there are j "switches" between dimensions, only states in $U^m \cup \cdots \cup U^{m+j}$ are involved. Thus, if the degree of α is j, the evaluation of α in $v \in U^m$ may involve states in U^{m+1}, U^{m+2}, $\ldots$, U^{m+j}, but not states from U^{m+j+1}. This means that the satisfaction relationship between α and (M_{m+j+1}, v) is classical — it does not depend on any state from U^{m+j+1}. ■

Finally, we define a $\mathcal{M}_{MCTL}$ model for ϕ_0. Let $j = \deg(\phi_0)$. We have that $M_{j+1}, \hat{s} \models \phi_0$ holds, and since $\hat{s} \in U^0$, the satisfaction relationship between ϕ_0 and $(M_{j+1}, \hat{s})$ is classical — the fact that $M_{j+1}, \hat{s} \models \phi_0$ does not depend on U^{j+1} (or τ^{j+1} or K^{j+1}). Take $M = (S, \{R_\delta : \delta \in \Delta\}, L)$ such that $S = S^{j+1}$, $R_\delta = R_\delta^{j+1}$, and $L(s) = L^{j+1}(s)$. Because $M_{j+1}, \hat{s} \models \phi_0$ does not depend on U^{j+1}, we also have that $M, \hat{s} \models \phi_0$. Since M_{j+1} is finite, M is finite.

4.3. Complexity

Now, we know that the satisfiability problem for CTL is EXPTIME-complete, and this gives us an EXPTIME lower bound for MCTL satisfiability (since CTL is — very obviously — a fragment of MCTL). But the construction described above also gives us an EXPTIME upper bound, thus giving the following.

THEOREM 4.11. *The satisfiability problem for* MCTL *is* EXPTIME-*complete.*

Consider the decision procedure described in Section 4.2. The idea behind this procedure is to use a constructive CTL decision procedure (such as the tableau method described in [8]) as a sub-routine for constructing components of a model for the input formula, each component corresponding to a different dimension. The use of the sub-routine is analytic, in that, each time we call the CTL satisfiability checking sub-routine, we are working with strict sub-formulae of the input formula. Thus, the overall running time of the procedure described in Section 4.2 for a formula ϕ over Θ is $O(2^{l \cdot m \cdot n})$ where $l = |dim(\phi)|$ is the number of dimensions in ϕ, $m = |\Theta|$ is the number of atomic propositions in ϕ, and $n = deg(\phi)$ is the degree of ϕ.

5. Normative Systems

Normative systems have come to play a major role in multi-agent systems research; for example, under the name of *social laws*, they have been shown to be a useful mechanism for coordination [21]. Following [22] and [1], for our purposes, a normative system is understood simply as *a set of constraints on the behaviour of agents in a system*. More precisely, a normative system defines, for every possible system transition, whether or not that transition is considered to be legal or not, in the context of the normative system. Different normative systems may differ on whether or not a particular transition is considered legal.

We now describe three related Normative Temporal Logics, NTL$^-$, NTL and NTL$^+$, the first of which is only a notational variant of MCTL. Rather than talking about a set of indices Δ we now assume a set of *norms* ∇. Given a set of atoms Θ and norms ∇, the language $\mathcal{L}_{NTL^-}(\Theta, \nabla)$ is defined by the following grammar ($p \in \Theta, \eta \in \nabla$):

$$\phi ::= \top \mid p \mid \neg\phi \mid \phi \vee \phi \mid \mathsf{P}^\eta \bigcirc \phi \mid \mathsf{P}^\eta(\phi \,\mathcal{U}\, \phi) \mid \mathsf{O}^\eta \bigcirc \phi \mid \mathsf{O}^\eta(\phi \,\mathcal{U}\, \phi)$$

The operators relate to the ones in the previous sections as follows: $\mathsf{P}^\eta\alpha$ equals $\mathsf{E}^\eta\alpha$, and $\mathsf{O}^\eta\alpha$ is $\mathsf{A}^\eta\alpha$. The intended meaning of $\mathsf{P}^\eta\alpha$ is 'given the norm η, α is *permitted*'. Likewise, $\mathsf{O}^\eta\alpha$ means that 'given the norm η, α is *obligatory*'.

We interpret norms in the context of a system: a system here will be simply a CTL model $M = (S, R, L)$, where R denotes all the transitions that the system could a-priori take. Now the interpretation $I(\eta)$ of a norm η w.r.t. M is simply a subset of R. These are the transitions that are *forbidden* by the norm, hence $R_\eta = R \backslash I(\eta)$ are those that are allowed, or *legal*. We require that R_η is total, expressing a *reasonabless* constraint: no norm prevents the system from making any further progress. Given a system $M = (S, R, L)$ and a set of norm symbols $\nabla = \{\eta, \ldots, \}$ a *system of norms* (based on M) is $\Upsilon = \langle S, R, L, I\rangle$, where the *interpretation* $I : \nabla \to 2^R$ is such that $R_\eta = R \backslash I(\eta)$ is a total relation. Observe that we can view systems of norms as MCTL models over dimensions $\Delta = \nabla$: take $R_\eta = R \setminus I(\eta)$. This defines the meaning $\Upsilon, s \models \varphi$ of a NTL$^-$ formula φ in the context of a system of norms and a state. Furthermore, the relationship between systems of norms and MCTL models (with ∇ as dimensions) is one-to-one, which immediately gives us a sound and complete axiomatisation $\mathcal{S}_{NTL^-}$ of NTL$^-$ as simply a notational variant of $\mathcal{S}_{MCTL}$. Moving on to NTL, we assume ∇ contains a symbol $\eta_\emptyset$ denoting the empty norm, and require that $I(\eta_\emptyset) = \emptyset$, for all I. The language $\mathcal{L}_{NTL}(\Theta, \nabla)$ extends that of $\mathcal{L}_{NTL^-}(\Theta, \nabla)$ with the symbol

$\eta_\emptyset$ which gives rise to the logic NTL. We will also write $\mathsf{A}\alpha$ for $\mathsf{O}^{\eta_\emptyset}\alpha$: this indicates what is naturally, or physically, inevitable in the system. Similarly, $\mathsf{E}\alpha = \mathsf{P}^{\eta_\emptyset}\alpha$ means that, when no restriction is imposed on the system, α is true along a path. However, in a formula, note that assumptions about the norms that are in place can be overruled by new occurrences of path operators: for instance, $\mathsf{P}^{\eta_1} \bigcirc \mathsf{P}^{\eta_2} \bigcirc \varphi$ means that it is possible to obey the norm η_1 for one transition, and then η_2 for the next one, and end up in a state where φ holds.

We have the following chain of implications in NTL. If something is naturally, or physically inevitable, then it is obligatory in any normative system; if something is an obligation within a given normative system η, then it is permissible in η; and if something is permissible in a given normative system, then it is naturally (physically) possible:

$$\models (\mathsf{A}\alpha \rightarrow \mathsf{O}^{\eta}\alpha) \qquad \models (\mathsf{O}^{\eta}\alpha \rightarrow \mathsf{P}^{\eta}\alpha) \qquad \models (\mathsf{P}^{\eta}\alpha \rightarrow \mathsf{E}\alpha)$$

The axioms $\mathcal{S}_{NTL}$ are those of $\mathcal{S}_{NTL^-}$ plus **Obl**: $\mathsf{O}^{\eta_\emptyset}\alpha \rightarrow \mathsf{O}^{\eta}\alpha$ and **Perm**: $\mathsf{P}^{\eta}\alpha \rightarrow \mathsf{P}^{\eta_\emptyset}\alpha$. Those axioms say that what is inevitable also holds after imposing any norm; that what is obligatory under a norm is permitted under that norm, and anything that is permitted is possible.

To show completeness of NTL, the same construction is used as for MCTL, treating $\eta_\emptyset$ as any other dimension, with the following difference. When expanding a node along dimension δ, when gluing the CTL model to the expanded node, label the transitions with $\eta_\emptyset$ in addition to δ. Axioms **Obl** and **Perm** ensure that this is consistent with the $\eta_\emptyset$-atoms present at the node.

EXAMPLE 5.1. Consider two parallel circular train tracks. At one point both tracks go through the same tunnel. At the east and the west end of the tunnel there are traffic lights, which can be either green or red. A train controller controls the lights. The eastern light should be set to green if and only if there is a train waiting to enter the east end of the tunnel and there is no train waiting at the west end of the tunnel, and similarly for the western light. One train travels on each of the tracks, in opposite directions. We call the train that enters the tunnel at the eastern end the east train and the other train the west train. Obviously, the trains should not enter the tunnel if the light is red.

We model this situation by considering the following system of norms. We assume that each train can be in one of three states: *tunnel* (the train is in the tunnel); *waiting* (the train is waiting to enter the tunnel); *away* (the train is neither in the tunnel nor waiting). When away, the train can either be away or waiting in the next state; when waiting the train can either

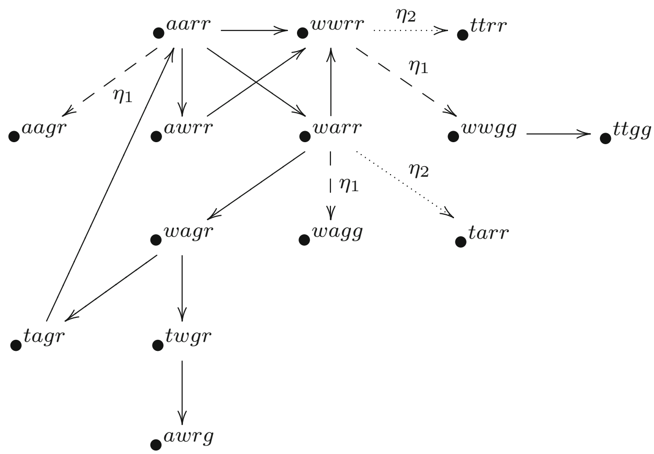

Figure 3. Υ for the trains example, including all physically possible transitions. Only part of the system is shown. The transitions prohibited by the normative systems η_1 and η_2 are shown with dashed and dotted lines, respectively. The labelling of the states is abbreviated for readability: "twgr" stands for tunnel-waiting-green-red and means that *wTunnel*, *eWaiting*, *wGreen* are true and that all other atoms (e.g., *eGreen*) are false.

be waiting or in the tunnel in the next state; when the train is in the tunnel it leaves the tunnel and is away in the next state. Thus, we use propositional atoms $eTunnel, eWaiting, eAway, wTunnel, wWaiting, wAway$ to encode the position of the east and west train. We also use atoms *eGreen* and *wGreen* to represent the fact that the eastern/western lights are green. Thus, $\neg eGreen$ means that the eastern light is red, and so on. Let $M = (S, R, L)$ be the Kripke structure where the states correspond to all possible configurations of the atomic propositions, let s be the state where both lights are red and both trains away, and the transitions are all physically possible transitions — illustrated in Figure 3. The transitions include entering on a red light, but exclude physically impossible transitions such as a train going directly from the *tunnel* state to the *waiting* state.

Let η_1 be the norm corresponding to the normative requirement on the switching of the lights described above: η_1 contains all transitions between states s_1 and s_2 in which one of the lights are set to green (in s_2) *without* the appropriate condition (as explained above) being true in s_1. The normative system η_1 is illustrated by labels on the transitions in Figure 3. The description above contains another normative requirement as well: trains should only enter the tunnel on a green light. Let η_2 be the normative system corresponding to that requirement: η_2 contains all transitions between states

s_1 and s_2 such that a train is in the tunnel in s_2 only if the corresponding light is green in s_1. Finally, let the norm η_3 be the combination of both norms: $\eta_3 = \eta_1 \cup \eta_2$. It is easy to see that $\nabla = \{\eta_1, \eta_2, \eta_3\}$, satisfies the reasonabless constraint in M. Let Υ be the system of norms using the interpretation of norms as just described.

While the norms in this particular example are designed to avoid a crash, there are other problems, such as "deadlock" (both trains can wait forever for a green light), which they do not avoid. For simplicity, we will only consider the norms mentioned above. Let the formula

$$crash = eTunnel \wedge wTunnel$$

denote a crash situation. We have that:

- $\Upsilon, s \models \mathsf{O}^{\eta_1} \bigcirc \neg wGreen$. In the initial state, according to normative system η_1 it is obligatory that the western light stays red in the next state.
- $\Upsilon, s \models \mathsf{P}^{\eta_1}(\neg eGreen\, \mathcal{U}\, eTunnel)$. η_1 permits the eastern light to stay red until the east train is in the tunnel.
- $\Upsilon, s \models \neg\mathsf{P}^{\eta_2}(\neg eGreen\, \mathcal{U}\, eTunnel)$. η_2 does not permit the eastern light to stay red until the east train is in the tunnel.
- $\Upsilon, s \models \mathsf{O}^{\eta_1} \Box(wGreen \rightarrow \neg eGreen)$. It is obligatory in the context of η_1 that at least one of the lights are red.
- $\Upsilon, s \models \mathsf{P}^{\eta_\emptyset} \Diamond crash$. Without any constraining norms, the system permits a crash in the future.
- $\Upsilon, s \models \mathsf{P}^{\eta_1} \Diamond crash$. The normative system η_1 permits a crash.
- $\Upsilon, s \models \mathsf{O}^{\eta_3} \Box \neg crash$. It is obligatory, in the context of η_3, that a crash never occurs; η_3 does not permit a crash at any point in the future.

It is worth reflecting on the compositional meaning of nested operators. For example, $\mathsf{P}^{\eta_3} \Diamond \mathsf{P}^{\eta_1} \bigcirc crash$ means that η_3 permits a computation along which in some future state $\mathsf{P}^{\eta_1} \bigcirc crash$ is true. However, in the evaluation of $\mathsf{P}^{\eta_1} \bigcirc crash$ in states along that computation, the system is *not* restricted by η_3 (but only by η_1). The following are examples of expressions involving nested operators.

- $\Upsilon, s \models \mathsf{O}^{\eta_\emptyset} \Box((wWaiting \wedge \neg wGreen) \rightarrow \neg \mathsf{P}^{\eta_2} \bigcirc wTunnel)$. It is obligatory in the system that it is always the case that if the west train is waiting and the western light is red then the western train is not permitted by η_2 in the tunnel in the next state.

- $\Upsilon, s \models \mathsf{P}^{\eta_2} \Diamond \mathsf{P}^{\eta_3} \bigcirc crash$. η_2 permits a future state where a crash in the next state is permitted even by η_3.
- $\Upsilon, s \models \mathsf{P}^{\eta_3} \Diamond \mathsf{P}^{\eta_1} \bigcirc crash$. η_3 permits a future state where a crash in the next state is permitted by η_1.
- $\Upsilon, s \models \mathsf{O}^{\eta_3} \Box \mathsf{O}^{\eta_2} \bigcirc \neg crash$. η_3 does not permit a future state where a crash is permitted in the next state by η_2.

Going beyond NTL, we can impose further structure on ∇ and its interpretations. For example, we can extend the logical language with basic statements such as $\eta \equiv \eta'$ and $\eta \sqsubset \eta'$ ($\sqsubseteq$ can then be defined), with the obvious interpretation.

Proposition 5.1. *Let $\Upsilon = \langle S, R, L, I \rangle$ be a system of norms, and $\eta_1, \eta_2 \in \nabla$. Then, if $I(\eta_1) \subseteq I(\eta_2)$ then $\Upsilon \models \mathsf{O}^{\eta_1}\phi \to \mathsf{O}^{\eta_2}\phi$ and $\Upsilon \models \mathsf{P}^{\eta_2}\phi \to \mathsf{P}^{\eta_1}\phi$.*

Furthermore, following [26], we can add unions and intersections of normative systems by requiring ∇ to include symbols $\eta \sqcup \eta'$, $\eta \sqcap \eta'$ whenever it includes η and η', and require interpretations to interpret $\sqcup$ as set union and $\sqcap$ as set intersection. Of course, we must then further restrict interpretations such that $R \setminus (I(\eta_1) \cup I(\eta_2))$ always is total. This would give us a kind of calculus of normative systems. Let Υ be a normative system with I being an interpretation with the mentioned properties:

i	$\Upsilon \models \mathsf{P}^{\eta \sqcup \eta'}\phi \to \mathsf{P}^{\eta}\phi$	*ii*	$\Upsilon \models \mathsf{P}^{\eta}\phi \to \mathsf{P}^{\eta \sqcap \eta'}\phi$
iii	$\Upsilon \models \mathsf{O}^{\eta}\phi \to \mathsf{O}^{\eta \sqcup \eta'}\phi$	*iv*	$\Upsilon \models \mathsf{O}^{\eta \sqcap \eta}\phi \to \mathsf{O}^{\eta}\phi$

These properties immediately follow from Proposition 5.1. They moreover seem rather natural: *i* for instance says that everything that is permitted while obeying two norms, is also permitted only obeying one of them. Having such a calculus allows one to reason about the composition of normative systems, similar to the way one constructs complex programs from simpler ones in Dynamic Logic [16].

We could drop the reasonableness constraint, making it possible that “too many” norms (i.e., too many constraints on agent behaviour) may prevent *any* transition from a given state. And many more research questions present themselves in the context of NTL: for instance, one may assign norms to agents and study what happens if some of them comply with their norms, while others don’t (cf. [4]). Similarly, one might consider prioritised collections of normative systems (“if that norm fails, then use this”).

6. Discussion

Model construction techniques similar to the one we have used are found in several works on transfer of properties to fusions. As discussed by Fajardo and Finger [10], many proofs of meta-logical properties of fusions in the literature [12, 13] employ the same strategy of (i) studying the modalisation/temporalisation of a generic logic; (ii) studying the (finite) *iterated modalisations* of two modal logics and (iii) viewing the fusion as a union of iterated modalisations. While the proof strategy used in this paper do not employ that strategy directly, there certainly are similarities.

The work presented in this paper has its roots in several different communities, the most significant being the tradition of using deontic logic in computer science to reason about normative behaviour of systems [24, 18], and the use of model checking and temporal logics such as CTL to analyse the temporal properties of systems [8, 6].

The two main differences between the language of NTL and the language of conventional deontic logic in computer science to reason about normative behaviour of systems [24, 18] are, first, *contextual* deontic operators allowing a formula to refer to several different normative systems, and, second, the use of *temporal* operators. All deontic expressions in NTL refer to time: $\mathsf{P}^\eta \bigcirc \phi$ ("it is permissible in the context of η that ϕ is true at the next time point"); $\mathsf{O}^\eta \Box \phi$ ("it is obligatory in the context of η that ϕ always will be true"); etc. Conventional deontic logic contains no notion of time. In order to compare our temporal deontic statements with those of deontic logic we must take the temporal dimension to be implicit in the latter. Two of the perhaps most natural ways of doing that is to take "obligatory" ($\mathsf{O}\phi$) to mean "*always* obligatory" ($\mathsf{O}^\eta \Box \phi$), or "obligatory at the *next point in time*" ($\mathsf{O}^\eta \bigcirc \phi$), respectively, and similarly for permission. In either case, all the principles of *Standard Deontic Logic* hold also for NTL, viz., $\mathsf{O}(\phi \rightarrow \psi) \rightarrow (\mathsf{O}\phi \rightarrow \mathsf{O}\psi)$ (K); $\neg\mathsf{O}\bot$ (D); and from ϕ infer $\mathsf{O}\phi$ (N). The two mentioned temporal interpretations of the (crucial) deontic axiom D are (both NTL validities):

$$\neg\mathsf{O}^\eta \Box \bot \text{ and } \neg\mathsf{O}^\eta \bigcirc \bot$$

Contrary-to-Duty obligations are structures involving two obligations, where the second obligation "takes over" when the first is violated [19]. Logics that deal with this kind of obligation typical add actions, time, a default component or a notion of context (signalling that the primary obligation has been violated, and we are now entering a sub-ideal context) to their semantic machinery to deal with them [19]. NTL is already equipped with a temporal component, and it would certainly also be possible to label the transitions in our semantics with actions. However, given that we incorporate a suite of norms within one system, it seems NTL can also represent "sub-ideal" contexts. We leave a detailed comparison between existing temporal deontic logics and NTL for future works, as well as any investigation into the usefulness of NTL to model contrary-to-duty obligations.

It has been argued that "*deadlines* are important norms in most interactions between agents" [7, page 40], and this naturally suggests the need for a temporal component in reasoning about systems with norms. We should also mention work by Sergot and his collaborators on the use of variants of the C+ language for representing and reasoning about deontic systems [20, 4]. The nC+ language they develop can be understood as language for defining representations of Kripke models and normative systems. Their work emerges from a rather different community — reasoning about action and non-monotonic reasoning in artificial intelligence. Finally, NTL is similar in spirit to the *deontic interpreted systems* model of [18]. Perhaps the most obvious difference is that while we consider "bad transitions", Lomuscio and Sergot are concerned with "bad states".

The design and application of normative systems and social laws is a major area of research activity in the multi-agent systems community. If we are going to make use of such social laws, then it seems only appropriate that we develop formalisms that allow us to explicitly and directly reason about them. We see the key advantages of NTL as follows. First, the fact that the formalism is so closely related to CTL is likely to be an advantage from the point of view of comprehension and acceptance within the mainstream model checking/verification community. Second, the fact that the language has a clear computational interpretation means that it can be applied in a computational setting without any ambiguity of interpretation. Third, the clear identification of different normative systems within the language, and the ability to talk about these directly, represents a novel step forward. While NTL arguably lacks some of the nuances of more conventional deontic and deontic temporal logics, we believe these advantages imply that the language and the approach it embodies merit further research.

Acknowledgements. JAR thanks Jose Castillejo (JC2008-00337) and projects IEA (TIN2006-15662-C02-01) and AT (CONSOLIDER CSD2007-0022).

References

[1] Ågotnes, T., W. van der Hoek, J.A. Rodriguez-Aguilar, C. Sierra, and M. Wooldridge, *On the logic of normative systems*, *IJCAI 2007*, pp. 1175–1180.

[2] Ågotnes, T., W. van der Hoek, and M. Wooldridge, 'Completeness and Complexity of Multi-Modal CTL', in C. Areces and S. Demri (eds.), *Proceedings of the 5th Workshop on Methods for Modalities (M4M5)*, (ENTCS), vol. 231, 2009, pp. 259–275.

[3] Ågotnes, T., W. van der Hoek, J.A. Rodríguez-Aguilar, C. Sierra, and M. Wooldridge, 'A Temporal Logic of Normative Systems', in D. Makinson, J. Malinowski and H. Wansing (eds.), *Towards Mathematical Philosophy*, Springer, 2008, pp. 69–106.

[4] Artikis, A., M.J. Sergot, and J.V. Pitt, 'Specifying norm-governed computational societies', *ACM Trans. Comput. Log.*, 10:1, 2009.

[5] Blackburn, P., M. de Rijke and Y. Venema, "Modal Logic," Cambridge Tracts in Theoretical Computer Science 53, CUP, 2001.

[6] Clarke, E. M., O. Grumberg and D. A. Peled, "Model Checking," MIT Press, 2000.

[7] Dignum, F., J. Broersen, V. Dignum, and J.-J.Ch. Meyer, 'Meeting the deadline: Why, when and how', vol. 3228 of *LNAI*, Springer, 2004, pp. 30–40.

[8] Emerson, E. A., 'Temporal and modal logic', in J. van Leeuwen (ed.), *Handbook of Theoretical Computer Science*, Elsevier, 1990, pp. 996–1072.

[9] Emerson, E. A., and J. Y. Halpern, 'Decision procedures and expressiveness in the temporal logic of branching time', *Journal of Computer and System Sciences* 30, 1985.

[10] Fajardo, R., and M. Finger, *How not to combine modal logics*, in B. Prasad (ed.), *IICAI*, 2005, pp. 1629–1647.

[11] Fine, K., and G. Schurz, 'Transfer theorems for multimodal logics', in J. Copeland (ed.), *Logic and Reality: Essays on the Legacy of Arthur Prior*, OUP, Oxford, 1996, pp. 169–213.

[12] Finger, M. and D. M. Gabbay, 'Adding a temporal dimension to a logic system', *Journal of Logic, Language, and Information* 1: 203–233, 1992.

[13] Finger, M. and M. A. Weiss, 'The unrestricted combination of temporal logic systems', *Logic Journal of the IGPL* 10: 165–189, 2002.

[14] Gabbay, D. M., A. Kurucz, F. Wolter and M. Zakharyaschev, *Many-Dimensional Modal Logics: Theory and Applications*, Elsevier, 2003.

[15] Gottlob, G, and C. Koch, *Monadic queries over tree-structured data*, in *LICS2002*, pp. 189–202.

[16] Harel, D., D. Kozen, and J. Tiuryn, *Dynamic Logic*, MIT Press, Cambridge, MA, 2000.

[17] Kracht, M. and F. Wolter, 'Simulation and transfer results in modal logic — a survey', *Studia Logica* 59: 149–177, 1997.

[18] Lomuscio, A., and M. Sergot, 'Deontic interpreted systems', *Studia Logica* 75 (1): 63–92, 2003.

[19] Prakken, H., and M. Sergot, 'Contrary-to-duty obligations', *Studia Logica* 57 (1): 91–115, 1996.

[20] Sergot, M., and R. Craven, 'The deontic component of action language nC+', in L. Goble and J.-J. Ch. Meyer (eds.), *DEON2006*, LNAI 4048, 2006, pp. 222–237.

[21] Shoham, Y., and M. Tennenholtz, 'On social laws for artificial agent societies: Off-line design', in P. E. Agre, and S. J. Rosenschein (eds.), *Computational Theories of Interaction and Agency*, MIT Press, 1996, pp. 597–618.

[22] Van der Hoek, W., M. Roberts and M. Wooldridge, 'Social Laws in Alternating Time: Effectiveness, Feasibility, and Synthesis', *Synthese* 156 (1): 1–19, 2007.

[23] XML Path Language, W3C Recommendation. `http://www.w3.org/TR/xpath`.

[24] Wieringa, R.J., and J.-J. Ch. Meyer, 'Deontic logic in computer science', J.-J. Ch. Meyer, and R. J. Wieringa (eds.), *Deontic Logic in Computer Science — Normative System Specification*, John Wiley & Sons, 1993, pp. 17–40.

[25] Wooldridge, M., *An Introduction to Multiagent Systems,* John Wiley & Sons, 2002.

[26] Wooldridge, M., and W. van der Hoek, 'On obligations and normative ability: Towards a logical analysis of the social contract', *Journal of Applied Logic*, 4 (3-4): 396–420, 2005.

Thomas Ågotnes
Department of Information Science
and Media Studies; University of Bergen
and Bergen University College
Bergen, Norway
`thomas.agotnes@infomedia.uib.no`

Wiebe van der Hoek &
Michael Wooldridge
Compute Science Department
University of Liverpool
Liverpool, UK
`{wiebe,mjw}@csc.liv.ac.uk`

Juan A. Rodríguez-Aguilar & Carles Sierra
Artificial Intelligence Research Institute – IIIA
Spanish Council for Scientific Research – CSIC
Spain
`{jar,sierra}@iiia.csic.es`