

REASONING ABOUT BOUNDS IN WEIGHTED TRANSITION SYSTEMS *

MIKKEL HANSEN, KIM GULDSTRAND LARSEN, RADU MARDARE,
AND MATHIAS RUGGAARD PEDERSEN

Department of Computer Science, Aalborg University, Denmark
e-mail address: {mhan,kgl,mardare,mrp}@cs.aau.dk

ABSTRACT. We propose a way of reasoning about minimal and maximal values of the weights of transitions in a weighted transition system (WTS). This perspective induces a notion of bisimulation that is coarser than the classic bisimulation: it relates states that exhibit transitions to bisimulation classes with the weights within the same boundaries. We propose a customized modal logic that expresses these numeric boundaries for transition weights by means of particular modalities. We prove that our logic is invariant under the proposed notion of bisimulation. We show that the logic enjoys the finite model property and we identify a complete axiomatization for the logic. Last but not least, we use a tableau method to show that the satisfiability problem for the logic is decidable.

1. INTRODUCTION

Weighted transition systems (WTSs) are used to model concurrent and distributed systems in the case where some resources are involved, such as time, bandwidth, fuel, or energy consumption. Recently, the concept of a cyber-physical system (CPS), which considers the integration of computation and the physical world has become relevant in modeling various real-life situations. In these models, sensor feedback affects computation, and through machinery, computation can further affect physical processes. The quantitative nature of weighted transition systems is well-suited for the quantifiable inputs and sensor measurements of CPSs, but their rigidity makes them less well-suited for the uncertainty inherent in CPSs. In practice, there is often some uncertainty attached to the resource cost, whereas weights in a WTS are precise. Thus, the model may be too restrictive and unable to capture the uncertainties inherent in the domain that is being modeled.

In this paper, we attempt to remedy this shortcoming by introducing a modal logic for WTSs that allows for approximate reasoning by speaking about upper and lower bounds for the weights of the transitions. The logic has two types of modal operators that reason about the minimal and maximal weights on transitions, respectively. This allows reasoning

Key words and phrases: Weighted transition system, bisimulation, axiomatization, completeness, satisfiability, modal logic, finite model property.

* This paper is an extended version of [HLM⁺16].

about models where the quantitative information may be imprecise (*e.g.* due to imprecisions introduced when gathering real data), but where we can establish a lower and upper bound for transitions.

In order to provide the semantics for this logic, we use the set of possible transition weights from one state to a set of states as an abstraction of the actual transition weights. The logic is expressive enough to characterize WTSs up to a relaxed notion of weighted bisimilarity, where the classical conditions are replaced with conditions requiring that the minimal and maximal weights on transitions are matched.

In [Ési14], Zoltán Ésik also considered the issue of bisimulation for weighted transition systems, although in the more general setting of synchronization trees with weights in an arbitrary monoid or semiring. Synchronization trees arise by unfolding the transitions of a weighted transition system starting in some state which will become the root of the tree. Both Ésik’s and our notion of bisimilarity bears some resemblance to probabilistic bisimulation [LS91], by considering not only single transitions but transitions to equivalence classes of states. However, while we require that the upper and lower bounds of these transitions should match, the bisimilarity of Ésik requires that the sum of the transitions should be the same. This is motivated by the fact that the synchronization trees do not form a category which respects the additive structure of a semiring. However, as Ésik proves, if one takes the quotient with respect to his version of weighted bisimilarity, then the category one obtains does respect the additive structure. Thus, the semiring structure of the weights is of vital importance to Ésik’s work, but is an aspect that we have not considered in our work.

Our main contribution is a complete axiomatization of our logic, showing that any validity in this logic can be proved as a theorem from the axiomatic system. Completeness allows us to transform any validity checking problem into a theorem proving one that can be solved automatically by modern theorem provers, thus bridging the gap to the theorem proving community. The completeness proof adapts the classical filtration method, which allows one to construct a (canonical) model using maximal consistent sets of formulae. The main difficulty of adapting this method to our setting is that we must establish both lower and upper bounds for the transitions in this model. To achieve this result, we demonstrate that our logic enjoys the finite model property.

Our second significant contribution is a decision procedure for determining the satisfiability of formulae in our logic. This decision procedure makes use of the tableau method to construct a tableau for a given formula. If the constructed tableau is successful, then the formula is satisfiable, and a finite model for the formula can be generated from the tableau.

Related Work. Several logics have been proposed in the past to express properties of quantified (weighted, probabilistic or stochastic) systems. They typically use modalities indexed with real numbers to express properties such as “ φ holds with at least probability b ”, “we can reach a state satisfying φ with a cost at least r ”, etc.

In the context of weighted automata, weighted monadic second order logic has been introduced by Droste and Gastin [DG05] to capture the behaviour of weighted automata for commutative semirings. This work has been extended to many closely related systems [BDP16][DR06][DV06][Mei06][Fic11]. There has also been work on connecting weighted monadic second order logic with probabilistic CTL [BG09]. For weighted transition systems, weighted modal logic has been introduced by Larsen and Mardare [LM14] to reason about the

consumption of resources in such a system. This logic has been extended to handle recursion [LMX14b][LMX14a] as well as parallel composition and concurrency [LMX15b]. For both the original weighted modal logic and its concurrent extension, complete axiomatizations were developed. A weighted extension of the μ -calculus was introduced by Larsen et al. in [LMX15a], where a complete axiomatization for this extension was also given.

While our setting is that of weighted transition systems, our logic and the development of its theory has more in common with Markovian logic than with the previously mentioned work on weighted systems.

Markovian logic was introduced by Mardare et al. [MCL12][CLM11a] building on previous work on probability logics [Zho09][FH94][HM01]. Markovian logic reasons about probabilistic and stochastic systems using operators L_r and M_r which mean that a property hold with *at least* probability r or *at most* probability r , respectively. Much of the work on Markovian logic has focused on giving a complete axiomatization for the logic [KMP13], culminating in a Stone duality for Markov processes [KLMP13]. However, compositional aspects have been considered in [CLM11b], where also an axiomatization was given for Markovian logic with an operator for parallel composition.

While our logical syntax resembles that of Markovian logic, our semantics is different in the sense that we argue not about probabilities, but about an interval of possible weights. For instance, in the aforementioned logics we have a validity of type $\vdash \neg L_r \phi \rightarrow M_r \phi$ saying that the value of the transition from the current state to ϕ is either at least r or at most r ; on the other hand, in our logic the formula $\neg L_r \phi \wedge \neg M_r \phi$ might have a model since $L_r \phi$ and $M_r \phi$ express the fact that the lower cost of a transition to ϕ is at least r and the highest cost is at most r respectively.

Our completeness proof uses a technique similar to the one used for weighted modal logic [LM14] and Markovian logic [KMP13][MCL12][CLM11a]. It is however different from these related constructions since our axiomatization is finitary, while the aforementioned ones require infinitary proof rules. Our axiomatic systems are related to the ones mentioned above and the mathematical structures revealed by this work are also similar to the related ones. This suggest a natural extension towards a Stone duality result along the lines of [KLMP13], which we will consider in a future work.

Decidability results regarding satisfiability have also been given for some related logics, such as weighted modal logic [LMX18] and probabilistic versions of CTL and the μ -calculus [CK16]. However, the satisfiability problem is known to be undecidable for other related logics, in particular timed logics such as TCTL [ACD93] and timed modal logic [JLMX14]. This fact suggests that our logic is an interesting one which, despite its expressivity, remains decidable.

Our approach of considering upper and lower bounds is related to interval-based formalisms such as interval Markov chains (IMCs) [JL91] and interval weighted modal transition systems (WMTSs) [JLS12]. Much like our approach, IMCs consider upper and lower bounds on transitions in the probabilistic case. WMTSs add intervals of weights to individual transitions of modal transition systems, in which there can be both may- and must-transitions. A main focus of the work both on IMCs and WMTSs have been a process of refinement, making the intervals progressively smaller until an implementation is obtained. However, none of these works have explored the logical perspective up to the level of axiomatization or satisfiability results, which is the focus of our paper.

2. MODEL

The models addressed in this paper are weighted transition systems, in which transitions are labeled with numbers to specify the cost of the corresponding transition. In order to specify and reason about properties regarding imprecision, such as “the maximum cost of going to a safe state is 10” and “the minimum cost of going to a halting state is 5”, we will abstract away the individual transitions and only consider the minimum and maximum costs from a state to another. We will do this by constructing for any two states the set of weights that are allowed from one to the other.

First we recap the definition of a weighted transition system. Let \mathcal{AP} be a countable set of atomic propositions. A WTS is formally defined as follows:

Definition 2.1. A *weighted transition system (WTS)* is a tuple $\mathcal{M} = (S, \rightarrow, \ell)$, where

- S is a non-empty set of *states*,
- $\rightarrow \subseteq S \times \mathbb{R}_{\geq 0} \times S$ is the *transition relation*, and
- $\ell : S \rightarrow 2^{\mathcal{AP}}$ is a *labeling function* mapping to each state a set of atomic propositions.

Note that we impose no restrictions on the state space S ; it can be uncountable. We write $s \xrightarrow{r} t$ to mean that $(s, r, t) \in \rightarrow$. We will say that a WTS is *image-finite* if for any $s \in S$ there are only finitely many $t \in S$ such that $s \xrightarrow{r} t$ for some $r \in \mathbb{R}_{\geq 0}$.

When modeling cyber-physical systems, it is often unreasonable to expect one to know the exact weights for transitions. However, it is often the case that one has some bounds on the actual weights, *e.g.* one might know that the cost of taking some transition is between 5 and 25. In order to reason about these bounds, we abstract away the individual transitions, and instead consider the set of weights between a state and a set of states.

Definition 2.2. For an arbitrary WTS $\mathcal{M} = (S, \rightarrow, \ell)$, the function $\theta_{\mathcal{M}} : S \rightarrow (2^S \rightarrow 2^{\mathbb{R}_{\geq 0}})$ is defined for any state $s \in S$ and set of states $T \subseteq S$ as

$$\theta_{\mathcal{M}}(s)(T) = \{r \in \mathbb{R}_{\geq 0} \mid \exists t \in T \text{ such that } s \xrightarrow{r} t\}.$$

Thus $\theta_{\mathcal{M}}(s)(T)$ is the set of all possible weights of going from s to a state in T . We will sometimes refer to $\theta(s)(T)$ as the *image from s to T* or simply as an *image set*. In the rest of the paper, we will use the notation

$$\theta^-(s)(T) = \begin{cases} -\infty & \text{if } \theta(s)(T) = \emptyset \\ \inf \theta(s)(T) & \text{otherwise} \end{cases}$$

and

$$\theta^+(s)(T) = \begin{cases} \infty & \text{if } \theta(s)(T) = \emptyset \\ \sup \theta(s)(T) & \text{otherwise.} \end{cases}$$

Thus $\theta^-(s)(T)$ will be a lower bound on the weights from s to T and $\theta^+(s)(T)$ will be an upper bound.

Example 2.3. Figure 1 shows a simple model of a robot vacuum cleaner that can be in a waiting state, a cleaning state, or a charging state. This is an example of a cyber-physical system where the costs of transitions are necessarily imprecise. The time it takes to recharge the batteries depends on the condition of the batteries as well as that of the charger; the time it takes to clean the room depends on how dirty the room is, and how free the floor is from obstacles; and the time it takes to reach the charger depends on where in the room the robot is when it needs to be recharged. By constructing the image sets, we can abstract




FIGURE 1. A simple model of a robot vacuum cleaner.

away from the individual transitions. For example, we have $\theta(s_2)(\{s_1\}) = \{5, 10, 15\}$, so $\theta^-(s_2)(\{s_1\}) = 5$ and $\theta^+(s_2)(\{s_1\}) = 15$.

We will now establish some useful properties of image sets. In particular, the transition function is monotonic with respect to set inclusion, and union distributes over image sets as one might expect.

Lemma 2.4 (Monotonicity of θ). *Let $\mathcal{M} = (S, \rightarrow, \ell)$ be a WTS and let T_1 and T_2 be subsets of S . If $T_1 \subseteq T_2$, then $\theta(s)(T_1) \subseteq \theta(s)(T_2)$.*

Lemma 2.5. *Let $\mathcal{M} = (S, \rightarrow, \ell)$ be a WTS. For any $s \in S$ and $T_1, T_2 \subseteq S$, it holds that*

- (1) $\theta(s)(T_1 \cup T_2) = \theta(s)(T_1) \cup \theta(s)(T_2)$ and
- (2) $\theta(s)(T_1 \cap T_2) \subseteq \theta(s)(T_1) \cap \theta(s)(T_2)$.

As usual we would like some way of relating model states with equivalent behavior. To this end we define the notion of a bisimulation relation. The classical notion of a bisimulation relation for weighted transition systems [BvBW06], which we term weighted bisimulation, is defined as follows.

Definition 2.6. Given a WTS $\mathcal{M} = (S, \rightarrow, \ell)$, an equivalence relation $\mathcal{R} \subseteq S \times S$ on S is called a *weighted bisimulation relation* iff for all $s, t \in S$, $s\mathcal{R}t$ implies

- (Atomic harmony) $\ell(s) = \ell(t)$,
- (Zig) if $s \xrightarrow{r} s'$ then there exists $t' \in S$ such that $t \xrightarrow{r} t'$ and $s'\mathcal{R}t'$, and
- (Zag) if $t \xrightarrow{r} t'$ then there exists $s' \in S$ such that $s \xrightarrow{r} s'$ and $s'\mathcal{R}t'$.

We say that $s, t \in S$ are weighted bisimilar, written $s \sim_W t$, iff there exists a weighted bisimulation relation \mathcal{R} such that $s\mathcal{R}t$. Weighted bisimilarity, \sim_W , is the largest weighted bisimulation relation.

Since it is our goal to abstract away from the exact weights on the transitions, the bisimulation that we will now introduce does not impose the classical zig-zag conditions [BvBW06] of a bisimulation relation, but instead require that bounds be matched for any bisimulation class.

Definition 2.7. Given a WTS $\mathcal{M} = (S, \rightarrow, \ell)$, an equivalence relation $\mathcal{R} \subseteq S \times S$ on S is called a *generalized weighted bisimulation relation* iff for all $s, t \in S$, $s\mathcal{R}t$ implies

- (Atomic harmony) $\ell(s) = \ell(t)$,
- (Lower bound) $\theta^-(s)(T) = \theta^-(t)(T)$, and
- (Upper bound) $\theta^+(s)(T) = \theta^+(t)(T)$

for any \mathcal{R} -equivalence class $T \subseteq S$.

Given $s, t \in S$ we say that s and t are generalized weighted bisimilar, written $s \sim t$, iff there exists a generalized weighted bisimulation relation \mathcal{R} such that $s\mathcal{R}t$. We let \sim denote generalized weighted bisimilarity which is defined as

$$\sim = \bigcup \{ \mathcal{R} \mid \mathcal{R} \text{ is a generalized weighted bisimulation relation} \}.$$

We will now show that generalized weighted bisimilarity, \sim , is the largest generalized weighted bisimulation relation. To this end, we first need to show that \sim is an equivalence relation.

Lemma 2.8. *Generalized weighted bisimilarity, \sim , is an equivalence relation.*

Proof. In order to prove that generalized weighted bisimilarity is an equivalence relation, we have to show that it is reflexive, symmetric and transitive.

Reflexivity: Consider the identity relation

$$\mathcal{I} = \{(s, s) \mid s \in S \text{ for some WTS } \mathcal{M} = (S, \rightarrow, \ell)\}.$$

It is trivial to verify that \mathcal{I} is a generalized weighted bisimulation relation, and therefore $\mathcal{I} \subseteq \sim$.

Symmetry: Let $\mathcal{M} = (S, \rightarrow, \ell)$ be a WTS and $s, t \in S$ states such that $s \sim t$. Because $s \sim t$ there must exist a generalized weighted bisimulation relation \mathcal{R} such that $s\mathcal{R}t$. Since \mathcal{R} is an equivalence relation, we immediately get $t \sim s$.

Transitivity: Let $\mathcal{M} = (S, \rightarrow, \ell)$ be a WTS and $s, t, u \in S$ states such that $s \sim t$ and $t \sim u$. There must exist generalized weighted bisimulation relations \mathcal{R} and \mathcal{R}' such that $s\mathcal{R}t$ and $t\mathcal{R}'u$. Let $\mathcal{R}'' = (\mathcal{R} \cup \mathcal{R}')^+$ be the transitive closure of the union of \mathcal{R} and \mathcal{R}' . Since \mathcal{R} and \mathcal{R}' are both equivalence relations, $\mathcal{R} \cup \mathcal{R}'$ is reflexive and symmetric, and since the transitive closure of a symmetric and reflexive relation is symmetric and reflexive, we get that \mathcal{R}'' is an equivalence relation. We need to show that \mathcal{R}'' is a generalized weighted bisimulation relation. Atomic harmony is trivially satisfied.

Suppose that $\theta(u)(T'') \neq \emptyset$ for some $T'' \in S/\mathcal{R}''$ implying the existence of a state $u' \in T''$ such that $\theta(u)(\{u'\}) \neq \emptyset$, further implying the existence of an equivalence class $T' \in S/\mathcal{R}'$ such that $u' \in T'$ and thus $\theta(u)(T') \neq \emptyset$. $t\mathcal{R}'u$ implies $\theta(t)(T') \neq \emptyset$ which further implies the existence of a state $t' \in T'$ such that $\theta(t)(\{t'\}) \neq \emptyset$. There must exist an equivalence class $T \in S/\mathcal{R}$ such that $t' \in T$ implying $\theta(t)(T) \neq \emptyset$. Because $s\mathcal{R}t$ we must have $\theta(s)(T) \neq \emptyset$ implying the existence of a state $s' \in T$ such that $\theta(s)(\{s'\}) \neq \emptyset$. $s', t' \in T$ implies $s'\mathcal{R}t'$, $t', u' \in T'$ implies $t'\mathcal{R}'u'$, and therefore $s'\mathcal{R}''u'$ implying $s' \in T''$ which further implies $\theta(s)(T'') \neq \emptyset$. Therefore $\theta(u)(T'') \neq \emptyset$ implies $\theta(s)(T'') \neq \emptyset$ for all $T'' \in S/\mathcal{R}''$. Symmetric arguments show that $\theta(s)(T'') \neq \emptyset$ implies $\theta(u)(T'') \neq \emptyset$ for all $T'' \in S/\mathcal{R}''$, and therefore $\theta(s)(T'') = \emptyset$ if and only if $\theta(u)(T'') = \emptyset$ for all $T'' \in S/\mathcal{R}''$.

Suppose towards a contradiction that $\theta^-(s)(T'') \neq \theta^-(u)(T'')$ for some $T'' \in S/\mathcal{R}''$. We have two cases to consider, namely $\theta^-(s)(T'') < \theta^-(u)(T'')$ and $\theta^-(s)(T'') > \theta^-(u)(T'')$. If $\theta^-(s)(T'') < \theta^-(u)(T'')$ there must exist a rational number $q \in \mathbb{Q}$ such that $\theta^-(s)(T'') < q < \theta^-(u)(T'')$, implying the existence of a state $s' \in T''$ such that $\theta^-(s)(T'') \leq \theta^-(s)(\{s'\}) < q$. There must exist $T \in S/\mathcal{R}$ such that $s' \in T$ implying $\theta^-(s)(T) < q$. Because $s\mathcal{R}t$ we must have $\theta^-(s)(T) = \theta^-(t)(T)$ implying the existence of a state $t' \in T$ such that $\theta^-(t)(\{t'\}) < q$. There must exist $T' \in S/\mathcal{R}'$ such that $t' \in T'$ implying $\theta^-(t)(T') < q$. Because $t\mathcal{R}'u$ we must have $\theta^-(t)(T') = \theta^-(u)(T')$ implying the existence of a state $u' \in T'$ such that $\theta^-(u)(\{u'\}) < q$. $s', t' \in T$ implies

$s'\mathcal{R}t', t', u' \in T'$ implies $t'\mathcal{R}u'$, and therefore $s'\mathcal{R}''u'$, implying $u' \in T''$ and therefore $\theta^-(u)(T'') < q$, leading to a contradiction. Symmetric arguments show that also $\theta^-(s)(T'') > \theta^-(u)(T'')$ leads to a contradiction and therefore $\theta^-(s)(T) = \theta^-(u)(T)$ for any $T \in S/\mathcal{R}''$.

Similar arguments show that $\theta^+(s)(T) = \theta^+(u)(T)$ for any $T \in S/\mathcal{R}''$ thus showing that \mathcal{R}'' is a generalized weighted bisimulation relation implying $\mathcal{R}'' \subseteq \sim$ and therefore $s \sim t$ and $t \sim u$ implies $s \sim u$. \square

Having established that \sim is an equivalence relation, we will now show that it is indeed the largest generalized weighted bisimulation relation.

Theorem 2.9. *Generalized weighted bisimilarity, \sim , is the largest generalized weighted bisimulation relation.*

Proof. We first show that \sim is a generalized weighted bisimulation relation. By Lemma 2.8 we know that \sim is an equivalence relation. Let $\mathcal{M} = (S, \rightarrow, \ell)$ be a WTS and $s, t \in S$ states such that $s \sim t$. There must exist a generalized weighted bisimulation relation \mathcal{R} such that $s\mathcal{R}t$, which trivially verifies atomic harmony.

Suppose that $\theta(t)(T) \neq \emptyset$ for some $T \in S/\sim$, implying the existence of a state $t' \in T$ such that $\theta(t)(\{t'\}) \neq \emptyset$. There must exist an equivalence class $T' \in S/\mathcal{R}$ such that $t' \in T'$, which implies that $\theta(t)(T') \neq \emptyset$. Because $s\mathcal{R}t$ we must have $\theta(s)(T') \neq \emptyset$, implying the existence of a state $s' \in T'$ such that $\theta(s)(\{s'\}) \neq \emptyset$. Because $s', t' \in T'$ we must have $s'\mathcal{R}t'$ and hence $s' \sim t'$, so $s' \in T$ and thus $\theta(s)(T) \neq \emptyset$. Symmetric arguments show that $\theta(s)(T) \neq \emptyset$ implies $\theta(t)(T) \neq \emptyset$ and therefore $\theta(s)(T) = \emptyset$ if and only if $\theta(t)(T) = \emptyset$ for all $T \in S/\sim$.

Suppose $\theta^-(s)(T) \neq \theta^-(t)(T)$ for some $T \in S/\sim$. We have two cases to consider, namely $\theta^-(s)(T) < \theta^-(t)(T)$ and $\theta^-(s)(T) > \theta^-(t)(T)$. If $\theta^-(s)(T) < \theta^-(t)(T)$ there must exist a rational number $q \in \mathbb{Q}$ such that $\theta^-(s)(T) < q < \theta^-(t)(T)$, implying the existence of a state $s' \in T$ such that $\theta^-(s)(T) \leq \theta^-(s)(\{s'\}) < q$. There must exist $T' \in S/\mathcal{R}$ such that $s' \in T'$ and hence $\theta^-(s)(T') < q$. Because $s\mathcal{R}t$ we have $\theta^-(s)(T') = \theta^-(t)(T')$, which means that there exists a state $t' \in T'$ such that $\theta^-(t)(\{t'\}) < q$. $s', t' \in T'$ implies $s'\mathcal{R}t'$ which further implies $s' \sim t'$ and therefore $\theta^-(t)(T) < q$, leading to a contradiction. Symmetric arguments show that also $\theta^-(s)(T) > \theta^-(t)(T)$ leads to a contradiction, and therefore $\theta^-(s)(T) = \theta^-(t)(T)$ for all $T \in S/\sim$.


Similar arguments show that $\theta^+(s)(T) = \theta^+(t)(T)$ for any $T \in S/\sim$, thus showing that \sim is a generalized weighted bisimulation relation.

\sim was defined as the union of all generalized weighted bisimulation relations, so for any generalized weighted bisimulation relation \mathcal{R} we must have $\mathcal{R} \subseteq \sim$, and hence we conclude that \sim is the largest generalized weighted bisimulation relation. \square

In what follows, we will use bisimulation to mean generalized weighted bisimulation and bisimilarity to mean generalized weighted bisimilarity.

Example 2.10. Consider the WTS depicted in Figure 2. It is easy to see that $\{s', t'\}$ is a \sim -equivalence class, and in fact it is the only \sim -equivalence class with in-going transitions. Since $\theta^-(s)(\{s', t'\}) = \theta^-(t)(\{s', t'\}) = 1$ and $\theta^+(s)(\{s', t'\}) = \theta^+(t)(\{s', t'\}) = 3$ we must have $s \sim t$, but because $s \xrightarrow{2} s'$ and $t \xrightarrow{2} t'$ it cannot be the case that $s \sim_W t$.

The following lemma shows that if two states are weighted bisimilar, then their image sets match exactly for any weighted bisimulation class.

FIGURE 2. $s \sim t$ but $s \not\sim_W t$.

Lemma 2.11. *Let $\mathcal{M} = (S, \rightarrow, \ell)$ be a WTS and let $s, t \in S$. $s \sim_W t$ implies that $\theta(s)(T) = \theta(t)(T)$ for any \sim_W -equivalence class $T \subseteq S$.*

Proof. Assume $s \sim_W t$ and let $T \subseteq S$ be a \sim_W -equivalence class. If $r \in \theta(s)(T)$, then there exists some $s' \in T$ such that $s \xrightarrow{r} s'$. Because $s \sim_W t$, there must exist some $t' \in T$ such that $t \xrightarrow{r} t'$ and $s' \sim_W t'$. Since T is a \sim_W -equivalence class, this means that $r \in \theta(t)(T)$. A similar argument shows that if $r \in \theta(t)(T)$, then $r \in \theta(s)(T)$. \square

We can now show the following relationship between \sim and \sim_W .

Theorem 2.12. *Generalized weighted bisimilarity is coarser than weighted bisimilarity, i.e.*

$$\sim_W \subseteq \sim \quad \text{and} \quad \sim_W \neq \sim.$$

Proof. Assume that $s \sim_W t$. We have that $\ell(s) = \ell(t)$, and by Lemma 2.11, we have that $\theta(s)(T) = \theta(t)(T)$ for any \sim_W -equivalence class $T \subseteq S$. This implies in particular that $\theta^-(s)(T) = \theta^-(t)(T)$ and $\theta^+(s)(T) = \theta^+(t)(T)$. Hence \sim_W is a bisimulation relation.

By Example 2.10, the inclusion is strict. \square

This result is not surprising, as our bisimulation relation only looks at the extremes of the transition weights, whereas weighted bisimulation looks at all of the transition weights.

3. LOGIC

In this section we introduce a modal logic which is inspired by Markovian logic [MCL12]. Our aim is that our logic should be able to capture the notion of bisimilar states as presented in the previous section, and as such it must be able to reason about the lower and upper bounds on transition weights.

Definition 3.1. The formulae of the logic \mathcal{L} are induced by the abstract syntax

$$\mathcal{L} : \quad \varphi, \psi ::= p \mid \neg\varphi \mid \varphi \wedge \psi \mid L_r\varphi \mid M_r\varphi$$

where $r \in \mathbb{Q}_{\geq 0}$ is a non-negative rational number and $p \in \mathcal{AP}$ is an atomic proposition.

L_r and M_r are modal operators. An illustration of how L_r and M_r are interpreted can be seen in Figure 3. Intuitively, $L_r\varphi$ means that the cost of transitions to where φ holds is *at least* r (see Figure 3a), and $M_r\varphi$ means that the cost of transitions to where φ holds is *at most* r (see Figure 3b). We now give the precise semantics interpreted over WTSs.



FIGURE 3. The semantics of L_r and M_r . If $\mathcal{M}, s \models L_r \varphi$, then r is to the left of $\theta^-(s)(\llbracket \varphi \rrbracket)$, and if $\mathcal{M}, s \models M_r \varphi$, then r is to the right of $\theta^+(s)(\llbracket \varphi \rrbracket)$.

Definition 3.2. Given a WTS $\mathcal{M} = (S, \rightarrow, \ell)$, a state $s \in S$ and a formula $\varphi \in \mathcal{L}$, the satisfiability relation \models is defined inductively as

$$\begin{aligned}
 \mathcal{M}, s \models p & \quad \text{iff} \quad p \in \ell(s), \\
 \mathcal{M}, s \models \neg \varphi & \quad \text{iff} \quad \mathcal{M}, s \not\models \varphi, \\
 \mathcal{M}, s \models \varphi \wedge \psi & \quad \text{iff} \quad \mathcal{M}, s \models \varphi \text{ and } \mathcal{M}, s \models \psi, \\
 \mathcal{M}, s \models L_r \varphi & \quad \text{iff} \quad \theta^-(s)(\llbracket \varphi \rrbracket_{\mathcal{M}}) \geq r, \\
 \mathcal{M}, s \models M_r \varphi & \quad \text{iff} \quad \theta^+(s)(\llbracket \varphi \rrbracket_{\mathcal{M}}) \leq r,
 \end{aligned}$$

where $\llbracket \varphi \rrbracket_{\mathcal{M}} = \{s \in S \mid \mathcal{M}, s \models \varphi\}$ is the set of all states of \mathcal{M} having the property φ .

We will omit the subscript \mathcal{M} from $\llbracket \varphi \rrbracket_{\mathcal{M}}$ whenever the model is clear from the context. If $\mathcal{M}, s \models \varphi$ we say that \mathcal{M} is a model of φ . A formula is said to be *satisfiable* if it has at least one model. We say that φ is a *validity* and write $\models \varphi$ if $\neg \varphi$ is not satisfiable. In addition to the operators defined by the syntax of \mathcal{L} , we also have the derived operators such as \perp , \rightarrow , etc. defined in the usual way. A *literal* is a formula that is of the form p or $\neg p$ where $p \in \mathcal{AP}$.

The formula $L_0 \varphi$ has special significance in our logic, as this formula means that there exists some transition to where φ holds. In fact, it follows in a straightforward manner from the semantics that $\mathcal{M}, s \models L_0 \varphi$ if and only if $\theta(s)(\llbracket \varphi \rrbracket) \neq \emptyset$. We can therefore encode the usual box and diamond modalities in our logic in the following way.

$$\diamond \varphi = L_0 \varphi \quad \square \varphi = \neg \diamond \neg \varphi.$$

Notice also that in general, the following schemes *do not hold*.

$$\begin{aligned}
 L_r \varphi \wedge L_r \psi & \rightarrow L_r(\varphi \wedge \psi) \\
 M_r \varphi \wedge M_r \psi & \rightarrow M_r(\varphi \wedge \psi)
 \end{aligned}$$

The reason that they do not hold in general is that there may be no transition to where $\varphi \wedge \psi$ holds, i.e. $\neg L_0(\varphi \wedge \psi)$. If we assume $L_0(\varphi \wedge \psi)$, then both schemes hold, as we show in Lemma 4.1. Another thing to note about the logic is that the formulae $L_r \varphi$ and $L_r \neg \varphi$ can both hold in the same model. To see this, simply construct a state that has two transitions with weight $x \geq r$ to two different states, one where φ holds and one where φ does not hold.

Example 3.3. Consider again our model of a robot vacuum cleaner depicted in Figure 1. Perhaps we want a guarantee that it takes no more than one time unit to go from a waiting state to a charging state. This can be expressed by the formula $\mathbf{waiting} \rightarrow M_1 \mathbf{charging}$, but since we know the only waiting state in our model is s_1 this can be simplified to simply checking whether $\mathcal{M}, s_1 \models M_1 \mathbf{charging}$. We thus have to check that $\theta^+(s_1)(\llbracket \mathbf{charging} \rrbracket) \leq$




FIGURE 4. s and t satisfy the same logical formulae, but $s \not\sim t$.

1. We do this by constructing the image set $\theta(s_1)(\llbracket \text{charging} \rrbracket)$. Since $\llbracket \text{charging} \rrbracket = \{s_3\}$, we have $\theta(s_1)(\{s_3\}) = \{1, 2\}$. Hence $\theta^+(s_1)(\llbracket \text{charging} \rrbracket) = 2 \not\subseteq 1$, so $\mathcal{M}, s_1 \not\models M_1 \text{charging}$.

Lemma 3.4. *Let $\mathcal{M} = (S, \rightarrow, \ell)$ be an image-finite WTS and $s \in S$. Let $T \subseteq S$ be a set such that all elements of T satisfy exactly the same formulae, and furthermore for any $t \in T$ and $t' \notin T$, there exists a formula φ such that $t \models \varphi$ and $t' \not\models \varphi$. Then there exists a formula $\varphi \in \mathcal{L}$ such that $\theta(s)(T) = \theta(s)(\llbracket \varphi \rrbracket)$.*

Proof. The idea of the proof is to repeatedly use the observation that if $t' \notin T$, then there exists a formula φ such that $t' \not\models \varphi$ and $t \models \varphi$ for all $t \in T$. First pick some formula φ_1 such that $t \models \varphi_1$ for all $t \in T$. Then $T \subseteq \llbracket \varphi_1 \rrbracket$, so $\theta(s)(T) \subseteq \theta(s)(\llbracket \varphi_1 \rrbracket)$. If $\theta(s)(T) \subsetneq \theta(s)(\llbracket \varphi_1 \rrbracket)$, then there must exist some $t_1 \notin T$ such that $s \xrightarrow{r} t_1$ and $t_1 \models \varphi_1$. Since $t_1 \notin T$, there must exist some formula φ_2 such that $t_1 \not\models \varphi_2$ and $t \models \varphi_2$ for all $t \in T$. We then get $\theta(s)(T) \subseteq \theta(s)(\llbracket \varphi_1 \wedge \varphi_2 \rrbracket)$. Again, if $\theta(s)(T) \subsetneq \theta(s)(\llbracket \varphi_1 \wedge \varphi_2 \rrbracket)$, then there must exist some $t_2 \notin T$ such that $s \xrightarrow{r} t_2$ and $t_2 \models \varphi_2$. Since $t_2 \notin T$, there must exist some formula φ_3 such that $t_2 \not\models \varphi_3$ and $t \models \varphi_3$ for all $t \in T$. Since \mathcal{M} is image-finite, there can only be finitely many states $t_i \notin T$ with $s \xrightarrow{r} t_i$, so continuing in the same way, we will eventually get a formula $\varphi_1 \wedge \dots \wedge \varphi_n$ such that $\theta(s)(T) = \theta(s)(\llbracket \varphi_1 \wedge \dots \wedge \varphi_n \rrbracket)$. \square

Next we show that our logic \mathcal{L} is invariant under bisimulation, which is also known as the Hennessy-Milner property. In order to prove this result, we have to restrict our models to only those that are image-finite, as shown by the following example.

Example 3.5. Consider the WTS depicted in Figure 4 with state space $S = \mathbb{N} \cup \{\omega, s, t\}$ and $\ell(s') = \emptyset$ for all $s' \in S$. The transition relation is given by $\omega \xrightarrow{0} \omega$, $s \xrightarrow{2} \omega$, $t \xrightarrow{3} \omega$, and $n + 1 \xrightarrow{0} n$, $s \xrightarrow{1} n$, and $t \xrightarrow{1} n$ for all $n \in \mathbb{N}$.

Then we have that $s_1 \sim s_2$ if and only if $s_1 = s_2$, since any states in $\mathbb{N} \cup \{\omega\}$ can be distinguished by the number of steps they can take, and s and t can be distinguished by the fact that $\theta^-(s)(\{\omega\}) = 2 \neq 3 = \theta^-(t)(\{\omega\})$. However, s and t satisfy all the same formulae, since any formula that holds in ω will also hold in n for some $n \in \mathbb{N}$, and the weights on the transitions to ω will therefore be masked by the bounds 1 and 4, and hence any formula can not distinguish between s and t .

The proof strategy follows a classical pattern: The left to right direction is shown by induction on φ for $\varphi \in \mathcal{L}$. The right to left direction is shown by constructing a relation \mathcal{R} relating those states that satisfy the same formulae and showing that this relation is a bisimulation relation.

Theorem 3.6 (Bisimulation invariance). *For any WTS $\mathcal{M} = (S, \rightarrow, \ell)$ and states $s, t \in S$ it holds that*

$$s \sim t \text{ implies } [\forall \varphi \in \mathcal{L}. \mathcal{M}, s \models \varphi \text{ iff } \mathcal{M}, t \models \varphi].$$

Furthermore, if \mathcal{M} is image-finite, then it also holds that

$$[\forall \varphi \in \mathcal{L}. \mathcal{M}, s \models \varphi \text{ iff } \mathcal{M}, t \models \varphi] \text{ implies } s \sim t.$$

Proof. We first show that $s \sim t$ implies $\mathcal{M}, s \models \varphi$ if and only if $\mathcal{M}, t \models \varphi$ for all $\varphi \in \mathcal{L}$ by induction on φ . The Boolean cases are trivial. If $\varphi = L_r\psi$, then we have $\theta^-(s)(\llbracket \psi \rrbracket) \geq r$, which implies that $\theta^-(s)(\llbracket \psi \rrbracket) \neq -\infty$. Assume towards a contradiction that $\theta^-(t)(\llbracket \psi \rrbracket) < r$. It can not be the case that $\theta^-(t)(\llbracket \psi \rrbracket) = -\infty$, hence it follows that $\llbracket \psi \rrbracket$ and $\theta^-(t)(\llbracket \psi \rrbracket)$ are non-empty, so there must exist some element $t' \in \llbracket \psi \rrbracket$ such that $\theta^-(t)(\llbracket \psi \rrbracket) \leq \theta^-(t)(\{t'\}) < r$. Since \sim is an equivalence relation, there must exist some \sim -equivalence class T such that $t' \in T$. This means that $\{t'\} \subseteq T$, so that also $\theta^-(t)(T) \leq \theta^-(t)(\{t'\}) < r$. By the induction hypothesis we have that $T \subseteq \llbracket \psi \rrbracket$. Because $s \sim t$, we have that $\theta^-(s)(T) = \theta^-(t)(T) < r$, so by monotonicity we get $\theta^-(s)(\llbracket \psi \rrbracket) \leq \theta^-(s)(T) < r$, which is a contradiction. The M_r case is handled similarly.

For the reverse direction, assume that \mathcal{M} is image-finite. We have to show that if for all $\varphi \in \mathcal{L}$, $\mathcal{M}, s \models \varphi$ if and only if $\mathcal{M}, t \models \varphi$ then $s \sim t$. To this end, we define a relation \mathcal{R} on S as

$$\mathcal{R} = \{(s, t) \in S \times S \mid \forall \varphi \in \mathcal{L}. \mathcal{M}, s \models \varphi \text{ iff } \mathcal{M}, t \models \varphi\}.$$

\mathcal{R} is clearly an equivalence relation and $s\mathcal{R}t$.

It is clear that $\ell(s) = \ell(t)$. Next we show that $\theta^-(s)(T) = \theta^-(t)(T)$ and $\theta^+(s)(T) = \theta^+(t)(T)$ for any \mathcal{R} -equivalence class T . Let $T \subseteq S$ be an \mathcal{R} -equivalence class. We first show that $\theta(s)(T) = \emptyset$ if and only if $\theta(t)(T) = \emptyset$. Assume that $\theta(s)(T) = \emptyset$. By Lemma 3.4 there exists a formula φ such that $\theta(s)(T) = \theta(s)(\llbracket \varphi \rrbracket) = \emptyset$, and therefore $s \not\models L_0\varphi$. Now assume towards a contradiction that $\theta(t)(T) \neq \emptyset$. Since \mathcal{M} is image-finite, there must be a finite subset $T' \subseteq T$ such that $\theta(t)(T) = \theta(t)(T')$. By Lemma 2.5, we then get $\theta(t)(T) = \bigcup_{t' \in T'} \theta(t)(\{t'\}) \neq \emptyset$, from which it follows that there must be some $t' \in T'$ such that $\theta(t)(\{t'\}) \neq \emptyset$. Since $t' \in T$, we must have $t' \models \varphi$, and therefore $t \models L_0\varphi$, which contradicts the fact that $s\mathcal{R}t$ and $s \not\models L_0\varphi$.

Now assume that $\theta(s)(T) \neq \emptyset$ and $\theta(t)(T) \neq \emptyset$. We need to show that $\theta^-(s)(T) = \theta^-(t)(T)$ and $\theta^+(s)(T) = \theta^+(t)(T)$. We do this by contradiction, which gives us four

(A1):	$\vdash \neg L_0 \perp$	
(A2):	$\vdash L_{r+q}\varphi \rightarrow L_r\varphi$	if $q > 0$
(A2'):	$\vdash M_r\varphi \rightarrow M_{r+q}\varphi$	if $q > 0$
(A3):	$\vdash L_r\varphi \wedge L_q\psi \rightarrow L_{\min\{r,q\}}(\varphi \vee \psi)$	
(A3'):	$\vdash M_r\varphi \wedge M_q\psi \rightarrow M_{\max\{r,q\}}(\varphi \vee \psi)$	
(A4):	$\vdash L_r(\varphi \vee \psi) \rightarrow L_r\varphi \vee L_r\psi$	
(A5):	$\vdash \neg L_0\psi \rightarrow (L_r\varphi \rightarrow L_r(\varphi \vee \psi))$	
(A5'):	$\vdash \neg L_0\psi \rightarrow (M_r\varphi \rightarrow M_r(\varphi \vee \psi))$	
(A6):	$\vdash L_{r+q}\varphi \rightarrow \neg M_r\varphi$	if $q > 0$
(A7):	$\vdash M_r\varphi \rightarrow L_0\varphi$	
(R1):	$\vdash \varphi \rightarrow \psi \implies \vdash (L_r\psi \wedge L_0\varphi) \rightarrow L_r\varphi$	
(R1'):	$\vdash \varphi \rightarrow \psi \implies \vdash (M_r\psi \wedge L_0\varphi) \rightarrow M_r\varphi$	
(R2):	$\vdash \varphi \rightarrow \psi \implies \vdash L_0\varphi \rightarrow L_0\psi$	

TABLE 1. The axioms for our axiomatic system, where $\varphi, \psi \in \mathcal{L}$ and $q, r \in \mathbb{Q}$.

cases to consider: $\theta^-(s)(T) < \theta^-(t)(T)$, $\theta^-(s)(T) > \theta^-(t)(T)$, $\theta^+(s)(T) < \theta^+(t)(T)$, and $\theta^+(s)(T) > \theta^+(t)(T)$.

For the case of $\theta^-(s)(T) < \theta^-(t)(T)$, there exists $q \in \mathbb{Q}_{\geq 0}$ such that

$$\theta^-(s)(T) < q < \theta^-(t)(T).$$

By Lemma 3.4, there exists a formula φ such that $\theta^-(t)(T) = \theta^-(t)(\llbracket \varphi \rrbracket)$. Since $T \subseteq \llbracket \varphi \rrbracket$, we then obtain

$$\theta^-(s)(\llbracket \varphi \rrbracket) \leq \theta^-(s)(T) < q < \theta^-(t)(T) = \theta^-(t)(\llbracket \varphi \rrbracket),$$

which implies that $s \not\models L_q\varphi$ but $t \models L_q\varphi$, and thus we get a contradiction. The other cases are handled similarly. \square

4. METATHEORY

In this section we propose an axiomatization for our logic that we prove not only sound, but also complete with respect to the proposed semantics.

4.1. Axiomatic System. Let $r, s \in \mathbb{Q}_{\geq 0}$. Then the deducibility relation $\vdash \subseteq 2^{\mathcal{L}} \times \mathcal{L}$ is a classical conjunctive deducibility relation, and is defined as the smallest relation which satisfies the axioms of propositional logic in addition to the axioms given in Table 1. We will write $\vdash \varphi$ to mean $\emptyset \vdash \varphi$, and we say that a formula or a set of formulae is *consistent* if it can not derive \perp .

The axioms presented in Table 1 bear some resemblance to the axiomatic systems of [MCL12] and [CLM11a]. Notably, our axiom A2 is almost identical to A2 of these works and capture similar properties about the systems being studied, with the major difference being that we reason about transition weights whereas the aforementioned works reason about rates or probabilities of transitions. Also worth noting here is the similarity between the rule R1 of these works and R1 of our axiomatic system. A notable difference is that we do not have the additive properties of measures for disjoint sets (since we are not working with probability measures), as is captured by the axioms A3 and A4 of these works. Also,

in one of the axiomatizations of [MCL12], the axioms A2 and A2' are not axioms, but can be derived from the axioms.

Rules R2 and R3 of [MCL12] and [CLM11a] reflect the Archimedean property of rationals, and while similar axioms can be proven sound in our setting, these were not needed to show our completeness result. We suspect, however, that if we were to pursue strong completeness, infinitary axioms similar to these would be needed.

Axiom A1 captures the notion that since \perp is never satisfied, we can never take a transition to where \perp holds. Axiom A2 says that if we know some value is the lower bound for going to where φ holds, then any lower value is also a lower bound for going to where φ holds. Axiom A2' is the analogue for upper bounds. Axioms A3-A4 show how L_r and M_r distribute over conjunction and disjunction. The version of axiom A4 where L_r is replaced with M_r is also sound, but as we show in Lemma 4.1, it can be proven from the other axioms. Axioms A5 and A5' say that if it is not possible to take a transition to where ψ holds, then including the states where ψ holds does not change the bounds. Axioms A6 and A7 show the relationship between L_r and M_r . In particular, A6 ensures that all bounds are well-formed. Notice also that the contrapositive of axiom A2 and A7 together gives us that $\neg L_0\varphi$ implies $\neg L_r\varphi$ and $\neg M_r\varphi$ for any $r \in \mathbb{Q}_{\geq 0}$. The rules R1 and R1' give a sort of monotonicity for L_r and M_r , and rule R2 says that if ψ follows from φ , then if it is possible to take a transition to where φ holds, it is also possible to take a transition to where ψ holds.

We now show some of the theorems which can be deduced from the axioms. T1, T1', and T5 together complete the distributivity properties for conjunction and disjunction. T2 and T2' make precise the intuitively clear idea that if two formulae are equivalent, then their upper and lower bounds should also be the same. T3 extends axiom A1 to hold for any $r \geq 0$, and T4 then extends this to any φ which implies \perp .

Lemma 4.1. *From the axioms listed in Table 1 we can derive the following theorems:*

- (T1): $\vdash (L_r\varphi \wedge L_q\psi \wedge L_0(\varphi \wedge \psi)) \rightarrow L_{\max\{r,q\}}(\varphi \wedge \psi)$
- (T1'): $\vdash (M_r\varphi \wedge M_q\psi \wedge L_0(\varphi \wedge \psi)) \rightarrow M_{\min\{r,q\}}(\varphi \wedge \psi)$
- (T2): $\vdash \varphi \leftrightarrow \psi \implies \vdash L_r\varphi \leftrightarrow L_r\psi$
- (T2'): $\vdash \varphi \leftrightarrow \psi \implies \vdash M_r\varphi \leftrightarrow M_r\psi$
- (T3): $\vdash \neg L_r\perp, \quad r \geq 0$
- (T4): $\vdash \varphi \rightarrow \perp \implies \vdash \neg L_r\varphi, \quad r \geq 0$
- (T5): $\vdash M_r(\varphi \vee \psi) \rightarrow M_r\varphi \vee M_r\psi$

Proof.

T1: Rule R1 implies

$$\vdash \neg L_q(\varphi \wedge \psi) \rightarrow (\neg L_q\varphi \vee \neg L_0(\varphi \wedge \psi)),$$

so also

$$\vdash \neg L_q(\varphi \wedge \psi) \rightarrow (\neg L_q\varphi \vee \neg L_0(\varphi \wedge \psi) \vee \neg L_r\psi).$$

This is equivalent to

$$\vdash (L_r\varphi \wedge L_q\psi \wedge L_0(\varphi \wedge \psi)) \rightarrow L_q(\varphi \wedge \psi).$$

T1': Similar to T1.

T2: Suppose $\vdash \varphi \leftrightarrow \psi$. We have that $\vdash L_r\varphi \rightarrow L_0\varphi$ by A2 and $\vdash L_0\varphi \rightarrow L_0\psi$ by R2.

Hence $\vdash L_r\varphi \rightarrow (L_r\varphi \wedge L_0\psi)$, so $\vdash L_r\varphi \rightarrow L_r\psi$ by R1. A similar argument shows that $\vdash L_r\psi \rightarrow L_r\varphi$, so $\vdash L_r\varphi \leftrightarrow L_r\psi$.

T2': Similar to T2.

T3: From axiom A1 we know that $\vdash \neg L_0 \perp$ which, by the contrapositive of A2, implies $\vdash \neg L_r \perp$ for any $r > 0$.

T4: Suppose $\vdash \varphi \rightarrow \perp$. We know for any $\psi \in \mathcal{L}$ that $\vdash \perp \rightarrow \psi$ and therefore $\vdash \varphi \rightarrow \perp \implies \vdash \varphi \leftrightarrow \perp$. From A1 we know that $\vdash \neg L_0 \perp$ and from T3 that $\vdash \neg L_r \perp$ for any $r > 0$ implying, by T2, that $\vdash \neg L_r \varphi$ for any $r \geq 0$.

T5: By axiom A7 we get $\vdash M_r(\varphi \vee \psi) \rightarrow L_0(\varphi \vee \psi)$ and A4 gives $\vdash L_0(\varphi \vee \psi) \rightarrow L_0\varphi \vee L_0\psi$. Hence we get $\vdash M_r(\varphi \vee \psi) \rightarrow (M_r(\varphi \vee \psi) \wedge L_0\varphi) \vee (M_r(\varphi \vee \psi) \wedge L_0\psi)$. Since $\vdash \varphi \rightarrow (\varphi \vee \psi)$ and $\vdash \psi \rightarrow (\varphi \vee \psi)$, rule R1' then gives $\vdash M_r(\varphi \vee \psi) \rightarrow M_r\varphi \vee M_r\psi$. \square

Next we prove that our axioms are indeed sound.

Theorem 4.2 (Soundness).

$$\vdash \varphi \text{ implies } \models \varphi.$$

Proof. The soundness of each axiom is easy to show, and many of them use the distributive property from Lemma 2.5. Here we prove the soundness for a few of the more interesting axioms.

A3: Suppose $\mathcal{M}, s \models L_r\varphi \wedge L_q\psi$ implying that $\mathcal{M}, s \models L_r\varphi$ and $\mathcal{M}, s \models L_q\psi$, implying further that $\theta^-(s)(\llbracket \varphi \rrbracket) \geq r$ and $\theta^-(s)(\llbracket \psi \rrbracket) \geq q$.

By Lemma 2.5 we must have that

$$\theta(s)(\llbracket \varphi \vee \psi \rrbracket) = \theta(s)(\llbracket \varphi \rrbracket \cup \llbracket \psi \rrbracket) = \theta(s)(\llbracket \varphi \rrbracket) \cup \theta(s)(\llbracket \psi \rrbracket)$$

and because $\theta^-(s)(\llbracket \varphi \rrbracket) \geq r$ and $\theta^-(s)(\llbracket \psi \rrbracket) \geq q$ we must have

$$\theta^-(s)(\llbracket \varphi \vee \psi \rrbracket) = \inf \theta(s)(\llbracket \varphi \rrbracket) \cup \theta(s)(\llbracket \psi \rrbracket) \geq \min \{r, q\}$$

implying $\mathcal{M}, s \models L_{\min\{r, q\}}(\varphi \vee \psi)$.

A4: Suppose $\mathcal{M}, s \models L_r(\varphi \vee \psi)$ implying that

$$\theta^-(s)(\llbracket \varphi \vee \psi \rrbracket) = \inf \theta(s)(\llbracket \varphi \rrbracket) \cup \theta(s)(\llbracket \psi \rrbracket) \geq r.$$

This implies that at least one of $\theta(s)(\llbracket \varphi \rrbracket)$ and $\theta(s)(\llbracket \psi \rrbracket)$ is non-empty. If $\theta(s)(\llbracket \varphi \rrbracket) \neq \emptyset$, then $\theta^-(s)(\llbracket \varphi \rrbracket) \geq r$, and also if $\theta(s)(\llbracket \psi \rrbracket) \neq \emptyset$, then $\theta^-(s)(\llbracket \psi \rrbracket) \geq r$, so at least one of $\mathcal{M}, s \models L_r\varphi$ and $\mathcal{M}, s \models L_r\psi$ must hold. Hence $\mathcal{M}, s \models L_r\varphi \vee L_r\psi$.

A6: Suppose $\mathcal{M}, s \models L_{r+q}\varphi$ implying that

$$\theta^-(s)(\llbracket \varphi \rrbracket) = \inf \theta(s)(\llbracket \varphi \rrbracket) \geq r + q.$$

It is clear that $\inf \theta(s)(\llbracket \varphi \rrbracket) \leq \sup \theta(s)(\llbracket \varphi \rrbracket)$, so

$$\theta^+(s)(\llbracket \varphi \rrbracket) = \sup \theta(s)(\llbracket \varphi \rrbracket) \geq \inf \theta(s)(\llbracket \varphi \rrbracket) \geq r + q > r.$$

Therefore, it cannot be the case that $\mathcal{M}, s \models M_r\varphi$ and thus $\mathcal{M}, s \models \neg M_r\varphi$.

R1: Suppose $\models \varphi \rightarrow \psi$ implying that $\llbracket \varphi \rrbracket \subseteq \llbracket \psi \rrbracket$, implying further, by the monotonicity of θ , that $\theta(s)(\llbracket \varphi \rrbracket) \subseteq \theta(s)(\llbracket \psi \rrbracket)$. Suppose further that $\mathcal{M}, s \models L_r\psi \wedge L_0\varphi$ implying $\mathcal{M}, s \models L_r\psi$ and $\mathcal{M}, s \models L_0\varphi$, implying further that

$$\theta^-(s)(\llbracket \psi \rrbracket) = \inf \theta(s)(\llbracket \psi \rrbracket) \geq r \quad \text{and} \quad \theta(s)(\llbracket \varphi \rrbracket) \neq \emptyset.$$

Since $\theta(s)(\llbracket \varphi \rrbracket)$ is non-empty, we then get that

$$\inf \theta(s)(\llbracket \varphi \rrbracket) \geq \inf \theta(s)(\llbracket \psi \rrbracket) \geq r,$$

which means that $\mathcal{M}, s \models L_r\varphi$. \square

4.2. Finite Model Property and Completeness. With our axiomatization proven sound we are now ready to present our main results, namely that our logic has the finite model property and that our axiomatization is complete.

To show the finite model property we will adapt the classical filtration method to our setting. Starting from an arbitrary formula ρ , we define a finite fragment of our logic, $\mathcal{L}[\rho]$, which we then use to construct a finite model for ρ . The main difference from the classical filtration method is that we must find an upper and a lower bound for the transitions in the model. For an arbitrary formula $\rho \in \mathcal{L}$ we define the following based on ρ :

- Let $Q_\rho \subseteq \mathbb{Q}_{\geq 0}$ be the set of all rational numbers $r \in \mathbb{Q}_{\geq 0}$ such that L_r or M_r appears in the syntax of ρ .
- Let Σ_ρ be the set of all atomic propositions $p \in \mathcal{AP}$ such that p appears in the syntax of ρ .
- The *granularity* of ρ , denoted as $gr(\rho)$, is the least common denominator of all the elements in Q_ρ .
- The *range* of ρ , denoted as R_ρ , is defined as

$$R_\rho = \begin{cases} \emptyset & \text{if } Q_\rho = \emptyset \\ I_\rho \cup \{0\} & \text{otherwise,} \end{cases}$$

where $I_\rho = \left\{ q \in \mathbb{Q}_{\geq 0} \mid \exists j \in \mathbb{N}. q = \frac{j}{gr(\rho)} \text{ and } \min Q_\rho \leq q \leq \max Q_\rho \right\}$. Here the granularity is used to pick out finitely many numbers in the interval. Note that we need to add 0 to R_ρ whether or not ρ actually contains 0 in any of its modalities. This is because, as we have pointed out before, formulae involving L_0 have special significance in our logic.

- The *modal depth* of ρ , denoted as $md(\rho)$, is defined inductively as:

$$md(\rho) = \begin{cases} 0 & \text{if } \rho = p \in \mathcal{AP} \\ md(\varphi) & \text{if } \rho = \neg\varphi \\ \max \{md(\varphi_1), md(\varphi_2)\} & \text{if } \rho = \varphi_1 \wedge \varphi_2 \\ 1 + md(\varphi) & \text{if } \rho = L_r\varphi \text{ or } \rho = M_r\varphi. \end{cases}$$

Since all formulae are finite, the modal depth is always a non-negative integer. The *language* of ρ , denoted by $\mathcal{L}[\rho]$, is defined as

$$\mathcal{L}[\rho] = \{\varphi \in \mathcal{L} \mid R_\varphi \subseteq R_\rho, md(\varphi) \leq md(\rho) \text{ and } \Sigma_\varphi \subseteq \Sigma_\rho\},$$

and we take $\mathcal{L}_{\leftrightarrow}[\rho]$ to be the Lindenbaum algebra of $\mathcal{L}[\rho]$, i.e. the quotient with respect to logical equivalence. The Lindenbaum algebra is a Boolean algebra with equivalence classes as elements. Note that the quotient $h : \mathcal{L}[\rho] \rightarrow \mathcal{L}_{\leftrightarrow}[\rho]$ is a homomorphism between Boolean algebras, and hence preserves the structure of $\mathcal{L}[\rho]$. For each element $x \in \mathcal{L}_{\leftrightarrow}[\rho]$, we fix now a formula $\varphi \in x$ to be the representative of that equivalence class, and we write $\hat{\varphi}$ for x . The order \leq in $\mathcal{L}_{\leftrightarrow}[\rho]$ is then given by $\hat{\varphi} \leq \hat{\psi}$ if and only if $\vdash \varphi \rightarrow \psi$. The join and meet in $\mathcal{L}_{\leftrightarrow}[\rho]$ are given by

$$\hat{\varphi} \vee \hat{\psi} = h(\varphi \vee \psi) \quad \hat{\varphi} \wedge \hat{\psi} = h(\varphi \wedge \psi),$$

and complement is given by

$$\neg \hat{\varphi} = h(\neg \varphi).$$

Note here the difference between $h(\varphi)$ and $\hat{\varphi}$. The quotient h sends φ to its equivalence class $x \in \mathcal{L}_{\leftrightarrow}[\rho]$. However, it may be the case that φ is not the representative for x , but some other formula ψ is. In that case we have $h(\varphi) = x = \hat{\psi}$. On the other hand, $\hat{\varphi}$ denotes

both that $\varphi \in \hat{\varphi}$, and also that φ is the chosen representative of its equivalence class, which ensures that in this case we have $h(\varphi) = \hat{\varphi}$.

The idea is that Σ_ρ ensures that only finitely many atomic propositions are used, R_ρ ensures that only finitely many weights on the modalities are used, and $md(\rho)$ puts a bound on the modal depth of formulae. The language $\mathcal{L}[\rho]$ itself is not finite, but contains only finitely many logically non-equivalent formulae. Hence $\mathcal{L}_{\leftrightarrow}[\rho]$ must be finite, and as we shall see, it contains all the information necessary to construct a model for ρ .

Proposition 4.3. *The language $\mathcal{L}_{\leftrightarrow}[\rho]$ is finite.*

Proof. Let $\mathcal{L}_{\leftrightarrow}^n[\rho]$ be the subset of $\mathcal{L}_{\leftrightarrow}[\rho]$ which only contains formulae of modal depth n . Then it is clear that

$$\mathcal{L}_{\leftrightarrow}[\rho] = \bigcup_{i=0}^{md(\rho)} \mathcal{L}_{\leftrightarrow}^i[\rho].$$

We will now prove by induction on the modal depth that for each i , $\mathcal{L}_{\leftrightarrow}^i[\rho]$ is finite.

$i = 0$: In this case, each element of $\mathcal{L}_{\leftrightarrow}^0[\rho]$ is a Boolean combination of atomic propositions in Σ_ρ . There are $2^{|\Sigma_\rho|}$ non-equivalent such formulae, so this set is finite.

$i > 0$: Each element of $\mathcal{L}_{\leftrightarrow}^i[\rho]$ is a Boolean combination of formulae of the form $L_r\varphi$ and $M_r\varphi$, where $\varphi \in \mathcal{L}_{\leftrightarrow}^j[\rho]$ for some $j < i$ and $r \in R_\rho$. By induction hypothesis, we know that there are only finitely many such φ . We know from Lemma 4.1 that if φ and ψ are logically equivalent, then $L_r\varphi$ and $L_r\psi$ as well as $M_r\varphi$ and $M_r\psi$ are also logically equivalent. Since R_ρ is finite, we conclude that $\mathcal{L}_{\leftrightarrow}^i[\rho]$ is finite. \square

In order to define the model, we need the standard notions of filters and ultrafilters on Boolean algebras [GH09]. A non-empty subset of a Boolean algebra B is called a *filter* if it is upward-closed with respect to the order, and closed under finite meets. A filter F is *proper* if $F \neq B$. An *ultrafilter* is a proper filter which is maximal in the sense of set inclusion.

The following property of ultrafilters is often useful.

Lemma 4.4. *For an ultrafilter F of $\mathcal{L}_{\leftrightarrow}[\rho]$ it holds that for any $\varphi \in \mathcal{L}[\rho]$, either $h(\varphi) \in F$ or $\neg h(\varphi) \in F$, but not both.*

We let $\mathcal{U}[\rho]$ denote the set of all ultrafilters on $\mathcal{L}_{\leftrightarrow}[\rho]$. Since $\mathcal{L}_{\leftrightarrow}[\rho]$ is finite, $\mathcal{U}[\rho]$ is also finite and consequently, any ultrafilter $u \in \mathcal{U}[\rho]$ must be a finite set. For any set $\Phi \subseteq \mathcal{L}_{\leftrightarrow}[\rho]$, the characteristic formula of Φ , denoted $\langle\!\langle \Phi \rangle\!\rangle$, is defined as

$$\langle\!\langle \Phi \rangle\!\rangle = \bigwedge_{\hat{\varphi} \in \Phi} \varphi.$$

Note that $\langle\!\langle \Phi \rangle\!\rangle \in \mathcal{L}[\rho]$ is a finite formula, and that if $u \in \mathcal{U}[\rho]$, then $h(\langle\!\langle u \rangle\!\rangle) \in u$.

We will now construct a (finite) model, \mathcal{M}_ρ , for ρ with state space $\mathcal{U}[\rho]$. In order to define the transition relation $\rightarrow_\rho \subseteq \mathcal{U}[\rho] \times \mathbb{R}_{\geq 0} \times \mathcal{U}[\rho]$, we consider any two ultrafilters $u, v \in \mathcal{U}[\rho]$ and define two functions $L, M : \mathcal{U}[\rho] \times \mathcal{U}[\rho] \rightarrow 2^{R_\rho}$ as

$$L(u, v) = \{r \mid h(L_r\langle\!\langle v \rangle\!\rangle) \in u\} \quad \text{and} \quad M(u, v) = \{s \mid h(M_s\langle\!\langle v \rangle\!\rangle) \in u\}.$$

The following lemma establishes a relationship between L and M , that we will need to define the transition relation. The lemma is a straightforward consequence of axiom A7.

Lemma 4.5. *Given any ultrafilters $u, v \in \mathcal{U}[\rho]$, it can not be the case that $L(u, v) = \emptyset$ and $M(u, v) \neq \emptyset$.*




FIGURE 5. When constructing a transition from u to v , we will only have information about what happens in the region Q_ρ and at 0. The line represents the non-negative real line and the arches represent the transitions that would be possible in a full model (i.e. one not restricted to $\mathcal{L}[\rho]$). The dashed part of the arches represent the part of the transition that we do not have information about.

Proof. Assume towards a contradiction that $L(u, v) = \emptyset$ and $M(u, v) \neq \emptyset$. Then we have $h(\neg L_0(v)) \in u$ and there exists some $r \in Q_\rho$ such that $h(M_r(v)) \in u$. However, by axiom A7, this implies that $h(L_0(v)) \in u$, which is a contradiction. \square

We can now define the transition relation in terms of $L(u, v)$ and $M(u, v)$. In Figure 5, we have illustrated the different cases that we must consider. Here, the area between $\min Q_\rho$ and $\max Q_\rho$ is the only part that the restricted language $\mathcal{L}[\rho]$ can speak about. The arches represent the interval within which transitions with that weight are possible. For any of the arches in the figure, we have the following correspondence with L_r and M_r .

- If a number r on the real line is contained within the arch, then we have $h(\neg L_r(v)) \in u$ and $h(\neg M_r(v)) \in u$.
- If a number r on the real line is to the left of the arch, then we have $h(L_r(v)) \in u$ and $h(\neg M_r(v)) \in u$.
- If a number r on the real line is to the right of the arch, then we have $h(M_r(v)) \in u$ and $h(\neg L_r(v)) \in u$.

In case (a) in Figure 5, we therefore have $L(u, v) \neq \emptyset$ and $M(u, v) \neq \emptyset$, so we have all the information we need to define the transition. In case (b) and (f), we have $L(u, v) \neq \emptyset$ and $M(u, v) = \emptyset$, since there exist numbers within the interval $[\min Q_\rho, \max Q_\rho]$ that are to the left of these arches, but none that are to the right. This means that we have enough information to define the minimum transition, but we do not know what the maximum transition is. Note that we can not simply say that the maximum transition is $\max Q_\rho$, because that would imply $h(M_{\max Q_\rho}(v)) \in u$, but we know that $M(u, v) = \emptyset$. Hence we need to pick a number that is to the right of $\max Q_\rho$ as the maximum. In case (d), we have both $L(u, v) = \emptyset$ and $M(u, v) = \emptyset$. This implies that $h(\neg L_0(v)) \in u$, which means that there should be no transition from u to v . In case (c) and (e), we have $L(u, v) = \emptyset$ and $M(u, v) \neq \emptyset$, but according to Lemma 4.5 these cases can never occur.

We therefore distinguish the following three cases in order to define the transition relation:

- (1) If $L(u, v) \neq \emptyset$ and $M(u, v) \neq \emptyset$, then we add the two transitions $u \xrightarrow{r_1} v$ and $u \xrightarrow{r_2} v$ where $r_1 = \max L(u, v)$ and $r_2 = \min M(u, v)$.
- (2) If $L(u, v) \neq \emptyset$ and $M(u, v) = \emptyset$, then we add the two transitions $u \xrightarrow{r_1} v$ and $u \xrightarrow{r_2} v$ where $r_1 = \max L(u, v)$ and $r_2 = \max Q_\rho + \frac{1}{gr(\rho)}$.
- (3) If $L(u, v) = \emptyset$ and $M(u, v) = \emptyset$, then there is no transition from u to v .

The following lemma tells us that these transitions are well-formed, i.e. that the lower bound on transitions is less than or equal to the upper bound.

Lemma 4.6. *For any ultrafilters $u, v \in \mathcal{U}[\rho]$, if $L(u, v) \neq \emptyset$ and $M(u, v) \neq \emptyset$, then $\max L(u, v) \leq \min M(u, v)$.*

Proof. Assume towards a contradiction that $\max L(u, v) > \min M(u, v)$. Then there exist $q, q' \in Q_\rho$ such that $q > q'$, $h(L_q(\uparrow v)) \in u$ and $h(M_{q'}(\uparrow v)) \in u$. Since $q > q'$, axiom A6 gives $h(\neg M_{q'}(\uparrow v)) \in u$, which is a contradiction. \square

Finally we define the labeling function $\ell_\rho : \mathcal{U}[\rho] \rightarrow 2^{\mathcal{AP}}$ for any $u \in \mathcal{U}[\rho]$ as $\ell_\rho(u) = \{p \in \mathcal{AP} \mid p \in u\}$. We then have a model $\mathcal{M}_\rho = (\mathcal{U}[\rho], \rightarrow_\rho, \ell_\rho)$, and it is not difficult to prove that \mathcal{M}_ρ is a WTS. Before we can prove the truth lemma, we need the following technical lemma.

Lemma 4.7. *For any consistent formula $\varphi \in \mathcal{L}[\rho]$, if $[\mathcal{M}_\rho, u \models \varphi \text{ iff } h(\varphi) \in u]$, then*

$$\bigvee_{v \in \llbracket \varphi \rrbracket} h(\uparrow v) \in u \quad \text{iff} \quad h(\varphi) \in u.$$

Proof. Suppose $\bigvee_{v \in \llbracket \varphi \rrbracket} h(\uparrow v) \in u$. Assume towards a contradiction that $h(\neg(\uparrow v)) \in u$ for all $v \in \llbracket \varphi \rrbracket$. Then, since u is an ultrafilter, we must have $\bigwedge_{v \in \llbracket \varphi \rrbracket} h(\neg(\uparrow v)) \in u$, which means that $\neg \bigvee_{v \in \llbracket \varphi \rrbracket} h(\uparrow v) \in u$, which is a contradiction. Hence there exists some $v' \in \llbracket \varphi \rrbracket$ such that $h(\uparrow v') \in u$. If $\hat{\psi} \in v'$, then $\vdash (\uparrow v') \rightarrow \psi$, so $\hat{\psi} \in u$ because u is an ultrafilter. Since $v' \in \llbracket \varphi \rrbracket$, we have by assumption that $h(\varphi) \in v'$, so we get $h(\varphi) \in u$.

Suppose $h(\varphi) \in u$, which by assumption means that $u \in \llbracket \varphi \rrbracket$, so $\vdash (\uparrow u) \rightarrow \bigvee_{v \in \llbracket \varphi \rrbracket} (\uparrow v)$. Since u is an ultrafilter, we have $h(\uparrow u) \in u$, and hence $\bigvee_{v \in \llbracket \varphi \rrbracket} h(\uparrow v) \in u$. \square

We are now in a position to state and prove the truth lemma, which says that an ultrafilter satisfies a formula in our model if and only if that formula is included in the ultrafilter.

Lemma 4.8 (Truth lemma). *If $\rho \in \mathcal{L}$ is a consistent formula, then for all $\varphi \in \mathcal{L}[\rho]$ and $u \in \mathcal{U}[\rho]$ we have*

$$\mathcal{M}_\rho, u \models \varphi \quad \text{iff} \quad h(\varphi) \in u.$$

Proof. The proof is by induction on the structure of φ . The Boolean cases are trivial. For the case $\varphi = L_r \psi$, we proceed as follows.

(\implies) Assume $\mathcal{M}_\rho, u \models L_r \psi$, meaning that $\theta^-(u)(\llbracket \psi \rrbracket) \geq r$. It can not be the case that $\theta^-(u)(\llbracket \psi \rrbracket) = \emptyset$, because otherwise $\theta^-(u)(\llbracket \psi \rrbracket) = -\infty$, and we have assumed $\theta^-(u)(\llbracket \psi \rrbracket) \geq r$. It also can not be the case that $\llbracket \psi \rrbracket = \emptyset$, because otherwise $\theta^-(u)(\llbracket \psi \rrbracket) = \emptyset$. We can partition all the ultrafilters $v \in \llbracket \psi \rrbracket$ as follows. Let $E = \{v \in \llbracket \psi \rrbracket \mid L(u, v) = \emptyset\}$ and $N = \{v \in \llbracket \psi \rrbracket \mid L(u, v) \neq \emptyset\}$. We then get that $E \cap N = \emptyset$, $E \cup N = \llbracket \psi \rrbracket$, $h(\neg L_0(\uparrow v)) \in u$ for

all $v \in E$, and $h(L_r(\ulcorner v \urcorner)) \in u$ for all $v \in N$. Because u is an ultrafilter, we then have

$$h\left(\bigwedge_{v \in E} \neg L_0(\ulcorner v \urcorner) \wedge \bigwedge_{v \in N} L_r(\ulcorner v \urcorner)\right) \in u.$$

By axiom A3, this implies

$$h\left(\bigwedge_{v \in E} \neg L_0(\ulcorner v \urcorner) \wedge L_r \bigvee_{v \in N} (\ulcorner v \urcorner)\right) \in u.$$

Then axiom A5 gives

$$h\left(L_r \bigvee_{v \in \llbracket \psi \rrbracket} (\ulcorner v \urcorner)\right) \in u.$$

By the induction hypothesis, T2, and Lemma 4.7, we then get $h(L_r \psi) \in u$.

(\Leftarrow) Let $h(L_r \psi) \in u$. It follows from A1, A2, and R2 that ψ is consistent. Hence, by the induction hypothesis, $\llbracket \psi \rrbracket$ is non-empty. We first show that $\theta(u)(\llbracket \psi \rrbracket) \neq \emptyset$. Assume therefore towards a contradiction that $\theta(u)(\llbracket \psi \rrbracket) = \emptyset$. Then for all $v \in \llbracket \psi \rrbracket$, we must have that case 3 holds, and hence $L(u, v) = \emptyset$, meaning $h(\neg L_r(\ulcorner v \urcorner)) \in u$ for all $v \in \llbracket \psi \rrbracket$. Since there are finitely many $v \in \llbracket \psi \rrbracket$, we can enumerate them as v_1, v_2, \dots, v_n . Then, since u is an ultrafilter, we have

$$h(\neg L_r(\ulcorner v_1 \urcorner) \wedge \neg L_r(\ulcorner v_2 \urcorner) \wedge \dots \wedge \neg L_r(\ulcorner v_n \urcorner)) \in u.$$

By De Morgan's law, this is equivalent to

$$h(\neg(L_r(\ulcorner v_1 \urcorner) \vee L_r(\ulcorner v_2 \urcorner) \vee \dots \vee L_r(\ulcorner v_n \urcorner))) \in u.$$

The contrapositive of axiom A4 then gives that

$$h(\neg L_r(\ulcorner v_1 \urcorner) \vee \ulcorner v_2 \urcorner \vee \dots \vee \ulcorner v_n \urcorner)) \in u,$$

and by the induction hypothesis, T2, and Lemma 4.7, this is equivalent to $\neg h(L_r \psi) \in u$, which is a contradiction.

Now assume towards a contradiction that $\theta^-(u)(\llbracket \psi \rrbracket) < r$. Then there exists some $v \in \llbracket \psi \rrbracket$ such that $\theta^-(u)(\ulcorner v \urcorner) < r$ and case 1 or case 2 holds. In either case we have $\max L(u, v) < r$ and hence there exists some $q \in Q_\rho$ such that $h(L_q(\ulcorner v \urcorner)) \in u$, which implies $h(L_0(\ulcorner v \urcorner)) \in u$ by axiom A2. By the induction hypothesis, $h(\psi) \in v$, which means that $\vdash \ulcorner v \urcorner \rightarrow \psi$. rule R1 then gives $h(L_r(\ulcorner v \urcorner)) \in u$, but this is a contradiction since $\max L(u, v) < r$.

The M_r case is similar, using axiom A7 instead of A2 to derive $h(L_0 \psi) \in u$. \square

Having established the truth lemma, we can now show that any consistent formula is satisfied by some finite model.

Theorem 4.9 (Finite model property). *For any consistent formula $\varphi \in \mathcal{L}$, there exists a finite WTS $\mathcal{M} = (S, \rightarrow, \ell)$ and a state $s \in S$ such that $\mathcal{M}, s \models \varphi$.*

Proof. Since $\varphi \in \mathcal{L}$ is consistent, $h(\varphi) \neq h(\perp)$, and since $\mathcal{L}_{\leftrightarrow}[\rho]$ is finite, there must exist an ultrafilter $u \in \mathcal{U}[\rho]$ such that $h(\varphi) \in u$. By the truth lemma, this means that $\mathcal{M}_\varphi, u \models \varphi$, and by construction, \mathcal{M}_φ is a finite model. \square

We are now able to state and prove our main result, namely that our axiomatization is complete.

Theorem 4.10 (Completeness). *For any formula $\varphi \in \mathcal{L}$, it holds that*

$$\models \varphi \text{ implies } \vdash \varphi.$$

Proof.

$$\models \varphi \text{ implies } \vdash \varphi$$

is equivalent to

$$\not\models \varphi \text{ implies } \not\vdash \varphi,$$

which is equivalent to

the consistency of $\neg\varphi$ implies the existence of a model for $\neg\varphi$,

and this is guaranteed by the finite model property. \square

We have thus established completeness for our logic. There is also a stronger notion of completeness, often called strong completeness, which asserts that $\Phi \models \varphi$ implies $\Phi \vdash \varphi$ for any set of formulae $\Phi \subseteq \mathcal{L}$. Completeness is a special case of strong completeness where $\Phi = \emptyset$. In the case of compact logics, strong completeness follows directly from completeness. However, our logic is non-compact.

Theorem 4.11. *Our logic is non-compact, meaning that there exists an infinite set $\Phi \subseteq \mathcal{L}$ such that each finite subset of Φ admits a model, but Φ does not.*

Proof. Consider the set $\Phi = \{L_q\varphi \mid q < r\} \cup \{\neg L_r\varphi\}$. For any finite subset of Φ , it is easy to construct a model. However, if $\mathcal{M}, s \models L_q\varphi$ for all $q < r$ where $q, r \in \mathbb{Q}_{\geq 0}$, then by the Archimedean property of the rationals, we also have $\mathcal{M}, s \models L_r\varphi$. Hence there can be no model for Φ . \square

5. SATISFIABILITY

The finite model property gives us a way of deciding in general whether there exists a model which satisfies a given formula. An algorithm would be to enumerate all finite models and all theorems derivable from the axioms, which can be done since there are countably many of each of these. If φ is satisfiable, it has a model, and by the finite model property, it has a finite one. So we can check one by one whether a finite model satisfies φ . On the other hand, if φ is not satisfiable, then $\neg\varphi$ is a theorem, so we can search through all theorems to see whether $\neg\varphi$ is one of them. Since φ is either satisfiable or its negation is a theorem, one of these two algorithms must eventually halt. By running these two algorithms in parallel, we have shown that the problem of deciding satisfiability for a given formula is decidable.

In what follows we do more: We propose an algorithm that constructs a tableau syntactically from a given formula. By inspecting this tableau, we can decide whether or not the formula is satisfiable, and if it is satisfiable, we can construct a model for the formula from the tableau.

As in the previous section, we impose an order on formulae given by $\varphi \leq \psi$ if and only if $\models \varphi \rightarrow \psi$. Given a finite set of formulae $\Gamma = \{\varphi_1, \dots, \varphi_n\}$, we denote by $\min(\Gamma)$ the set of minimal elements of Γ , i.e.

$$\min(\Gamma) = \{\varphi_i \in \Gamma \mid \text{there is no } \varphi_j \text{ such that } \varphi_j \leq \varphi_i\},$$

and we let

$$\mathcal{L}(\Gamma) = \{\varphi_i \in \Gamma \mid \text{there is no } j < i \text{ such that } \models \varphi_j \leftrightarrow \varphi_i\}.$$

$$\begin{array}{c}
(\wedge) \frac{\langle \Gamma \cup \{\varphi \wedge \psi\}, \mathcal{I}^L, \mathcal{I}^M \rangle}{\langle \Gamma \cup \{\varphi, \psi\}, \mathcal{I}^L, \mathcal{I}^M \rangle} \quad (\neg \wedge) \frac{\langle \Gamma \cup \{\neg(\varphi \wedge \psi)\}, \mathcal{I}^L, \mathcal{I}^M \rangle}{\langle \Gamma \cup \{\neg\varphi\}, \mathcal{I}^L, \mathcal{I}^M \rangle \quad \langle \Gamma \cup \{\neg\psi\}, \mathcal{I}^L, \mathcal{I}^M \rangle} \\
(\neg \neg) \frac{\langle \Gamma \cup \{\neg\neg\varphi\}, \mathcal{I}^L, \mathcal{I}^M \rangle}{\langle \Gamma \cup \{\varphi\}, \mathcal{I}^L, \mathcal{I}^M \rangle} \\
(\text{mod}) \frac{\langle \Gamma \cup \{N_{r_1}^1 \varphi_1, \dots, N_{r_n}^n \varphi_n\} \cup \{-O_{r'_1}^1 \varphi'_1, \dots, -O_{r'_n}^n \varphi'_n\}, \mathcal{I}^L, \mathcal{I}^M \rangle}{\langle \{\psi_1\}, \mathcal{I}_1^L, \mathcal{I}_1^M \rangle \quad \dots \quad \langle \{\psi_k\}, \mathcal{I}_k^L, \mathcal{I}_k^M \rangle} \\
\text{if } N^i \in \{L, M\} \text{ for all } 1 \leq i \leq n, O^j \in \{L, M\} \text{ for all } \\
1 \leq j \leq n', \text{ and no formula in } \Gamma \text{ is of the form } N_r \varphi \text{ or } \neg N_r \varphi \\
\text{where } N \in \{L, M\}.
\end{array}$$

TABLE 2. Tableau rules

Furthermore, we let $\uparrow_\Gamma(\varphi)$ be the upward closure of φ in Γ , i.e.

$$\uparrow_\Gamma(\varphi) = \{\varphi' \in \Gamma \mid \varphi \leq \varphi'\}.$$

A *tableau* is a tree with nodes of the form $\langle \Gamma, \mathcal{I}^L, \mathcal{I}^M \rangle$ that is constructed from the rules of Table 2, where the (mod) rule may only be used when no other rule can be used. For each node $\langle \Gamma, \mathcal{I}^L, \mathcal{I}^M \rangle$, Γ is a set of formulae, and \mathcal{I}^L and \mathcal{I}^M are intervals of the form $\lambda a, b \int$ where $a \in \mathbb{R}_{\geq 0} \cup \{-\infty\}$, $b \in \mathbb{R}_{\geq 0} \cup \{\infty\}$, $\lambda \in \{[, (, \text{ and } \int \in [],)\}$, subject to the constraint that $\lambda = ($ if $a = -\infty$ and $\int =)$ if $b = \infty$. We will say that an interval $\lambda a, b \int$ is *consistent* if $a < b$ or $a = b$ and the interval is closed.

For the rule (mod), the objects ψ_i , \mathcal{I}_i^L and \mathcal{I}_i^M in the conclusion are constructed as follows. The ψ_i are given by

$$\{\psi_1, \dots, \psi_k\} = \min(\mathcal{L}(\{\varphi_1, \dots, \varphi_n\})).$$

We will show later how to actually compute $\{\psi_1, \dots, \psi_k\}$. Let $\Gamma' = \{\varphi_1, \dots, \varphi_n\}$ and

$$\mathbb{L}_i^+ = \{r \mid L_r \varphi_j = N_{r_j}^j \varphi_j \text{ for some } j \text{ and } \varphi_j \in \uparrow_{\Gamma'}(\psi_i)\}$$

$$\mathbb{M}_i^+ = \{r \mid M_r \varphi_j = N_{r_j}^j \varphi_j \text{ for some } j \text{ and } \varphi_j \in \uparrow_{\Gamma'}(\psi_i)\}$$

as well as

$$\mathbb{L}_i^- = \{r \mid L_r \varphi'_j = O_{r_j}^j \varphi'_j \text{ for some } j \text{ and } \models \psi_i \rightarrow \varphi'_j\}$$

$$\mathbb{M}_i^- = \{r \mid M_r \varphi'_j = O_{r_j}^j \varphi'_j \text{ for some } j \text{ and } \models \psi_i \rightarrow \varphi'_j\}.$$

Then the intervals \mathcal{I}_i^L and \mathcal{I}_i^M are given by

$$\mathcal{I}_i^L = \begin{cases} [\max \mathbb{L}_i^+, \min \mathbb{L}_i^-) & \text{if } \mathbb{L}_i^+ \neq \emptyset \text{ and } \mathbb{L}_i^- \neq \emptyset \\ [0, \min \mathbb{L}_i^-) & \text{if } \mathbb{L}_i^+ = \emptyset \text{ and } \mathbb{L}_i^- \neq \emptyset \\ [\max \mathbb{L}_i^+, \infty) & \text{if } \mathbb{L}_i^+ \neq \emptyset \text{ and } \mathbb{L}_i^- = \emptyset \\ [0, \infty) & \text{if } \mathbb{L}_i^+ = \emptyset \text{ and } \mathbb{L}_i^- = \emptyset \end{cases}$$

$$\mathcal{I}_i^M = \begin{cases} (\max \mathbb{M}_i^-, \min \mathbb{M}_i^+] & \text{if } \mathbb{M}_i^- \neq \emptyset \text{ and } \mathbb{M}_i^+ \neq \emptyset \\ [0, \min \mathbb{M}_i^+] & \text{if } \mathbb{M}_i^- = \emptyset \text{ and } \mathbb{M}_i^+ \neq \emptyset \\ (\max \mathbb{M}_i^-, \infty) & \text{if } \mathbb{M}_i^- \neq \emptyset \text{ and } \mathbb{M}_i^+ = \emptyset \\ [0, \infty) & \text{if } \mathbb{M}_i^- = \emptyset \text{ and } \mathbb{M}_i^+ = \emptyset \end{cases}$$

Informally, one should think of a node $m = \langle \Gamma, \mathcal{I}^L, \mathcal{I}^M \rangle$ as satisfying all the formulas in Γ . Moreover, the (mod)-rule signifies a state transition, where the new states are given by the nodes in the conclusion, and any transition to m must have a minimum weight that lies in the interval \mathcal{I}^L , and a maximum weight that lies in the interval \mathcal{I}^M .

Example 5.1. We now illustrate the use of the (mod) rule through an example. Consider the node $m = \langle \{p_1, p_2, L_2p_1, L_4(p_1 \wedge p_2), L_0p_3, \neg L_5p_2, \neg M_6p_3\}, \mathcal{I}^L, \mathcal{I}^M \rangle$. We group the formulas as

$$\Gamma = \{p_1, p_2\}, \Gamma' = \{L_2p_1, L_4(p_1 \wedge p_2), L_0p_3\}, \text{ and } \Gamma'' = \{\neg L_5p_2, \neg M_6p_3\},$$

so that $m = \langle \Gamma \cup \Gamma' \cup \Gamma'', \mathcal{I}^L, \mathcal{I}^M \rangle$. Since Γ only includes literals, it is clear that we can use no other rules, so we are allowed to use (mod) on m .

We see that $\models (p_1 \wedge p_2) \rightarrow p_1$, and hence $\{\psi_1, \psi_2\} = \{p_1 \wedge p_2, p_3\}$, so there are two children of m . For the first child, we find

$$\begin{aligned} \mathbb{L}_1^+ &= \{2, 4\} & \mathbb{M}_1^+ &= \emptyset \\ \mathbb{L}_1^- &= \{5\} & \mathbb{M}_1^- &= \emptyset, \end{aligned}$$

and for the second child we find

$$\begin{aligned} \mathbb{L}_2^+ &= \{0\} & \mathbb{M}_2^+ &= \emptyset \\ \mathbb{L}_2^- &= \emptyset & \mathbb{M}_2^- &= \{6\}. \end{aligned}$$

Hence the intervals become

$$\begin{aligned} \mathcal{I}_1^L &= [4, 5) & \mathcal{I}_1^M &= [0, \infty) \\ \mathcal{I}_2^L &= [0, \infty) & \mathcal{I}_2^M &= (6, \infty), \end{aligned}$$

and our application of the rule becomes

$$\text{(mod)} \frac{\langle \{p_1, p_2, L_2p_1, L_4(p_1 \wedge p_2), L_0p_3, \neg L_5p_2, \neg M_6p_3\}, \mathcal{I}^L, \mathcal{I}^M \rangle}{\langle \{p_1 \wedge p_2\}, [4, 5), [0, \infty) \rangle \quad \langle \{p_3\}, [0, \infty)(6, \infty) \rangle}$$

Given a formula φ , we will say that a tableau \mathcal{T} is a *tableau for* φ if $\langle \{\varphi\}, [0, 0], [0, 0] \rangle$ is the root of \mathcal{T} .

Definition 5.2. A node m in a tableau is called

- a *modal node* if the (mod)-rule was applied to m and
- a *terminal node* if it is either a modal node or a leaf node.

Definition 5.3. A node $m = \langle \Gamma, \lceil_1 a, b \rceil_1, \lceil_2 c, d \rceil_2 \rangle$ is *consistent* if

- for any $p \in \mathcal{AP}$ we do not have both $p \in \Gamma$ and $\neg p \in \Gamma$,
- $\lceil_1 a, b \rceil_1$ and $\lceil_2 c, d \rceil_2$ are consistent, and
- either $a < d$ or $a = d$, $\lceil_1 = [$, and $\rceil_2 =]$.

Definition 5.4. A tableau \mathcal{T} is *successful* if there exists a subtree \mathcal{T}' of \mathcal{T} such that

- every leaf in \mathcal{T}' is also a leaf in \mathcal{T} ,
- if a modal node m is included in \mathcal{T}' , then every child of m is also included in \mathcal{T}' , and

- every terminal node in \mathcal{T}' is consistent.

Given a successful tableau \mathcal{T} , we construct the WTS $\mathcal{M}(\mathcal{T})$ with state $s_{\mathcal{T}}$ using Algorithm 1.

```

1 Let  $\mathcal{T}'$  be a witness for the fact that  $\mathcal{T}$  is successful ;
2  $S := \{s_{\mathcal{T}}\}, \rightarrow := \emptyset, \ell := \emptyset$  ;
3 Let  $X$  be a stack and  $X := \emptyset$  ;
4  $X.push((s_{\mathcal{T}}, r))$  where  $r$  is the root of  $\mathcal{T}'$  ;
5 while  $X \neq \emptyset$  do
6    $(s, m) := X.pop$  ;
7   Let  $m = \langle \Gamma, \Delta, (a, b) \rangle$  ;
8   if  $m$  is not a terminal node then
9     Let  $m'$  be the left-most child of  $m$  in  $\mathcal{T}'$  ;
10     $X.push((s, m'))$  ;
11  end
12  if  $m$  is a leaf node then
13     $\ell := \ell \cup \{(s, p) \mid p \in \mathcal{AP} \text{ and } p \in \Gamma\}$  ;
14  end
15  if  $m$  is a modal node then
16     $\ell := \ell \cup \{(s, p) \mid p \in \mathcal{AP} \text{ and } p \in \Gamma\}$  ;
17    Let  $m_1 = \langle \Gamma_1, \mathcal{I}_1^L, \mathcal{I}_1^M \rangle, \dots, m_n = \langle \Gamma_n, \mathcal{I}_n^L, \mathcal{I}_n^M \rangle$  be the children of  $m$  in  $\mathcal{T}'$  ;
18    for  $i = 1, \dots, n$  do
19      Let  $\mathcal{I}_i^L = \lceil a_i, b_i \rceil$  and  $\mathcal{I}_i^M = \lceil c_i, d_i \rceil$  ;
20       $x_i := a_i$  ;
21       $y_i := \begin{cases} \max\{a_i, \frac{d_i - c_i}{2} + c_i\} & \text{if } d_i \neq \infty \\ \max\{a_i, c_i + 1\} & \text{if } d_i = \infty \end{cases}$  ;
22       $S := S \cup \{s_i\}$  ;
23       $\rightarrow := \rightarrow \cup \{(s, x_i, s_i), (s, y_i, s_i)\}$  ;
24       $X.push((s_i, m_i))$  ;
25    end
26  end
27 end
28  $\mathcal{M}(\mathcal{T}) := (S, \rightarrow, \ell)$  ;
29 return  $(\mathcal{M}(\mathcal{T}), s_{\mathcal{T}})$  ;

```

Algorithm 1: Constructing the model $\mathcal{M}(\mathcal{T})$ for a successful tableau \mathcal{T} .

Lemma 5.5. *If \mathcal{T} is a successful tableau for φ , then $\mathcal{M}(\mathcal{T}), s_{\mathcal{T}} \models \varphi$.*

Proof. Let Y be the set of all pairs (s, m) that are added to the stack X by Algorithm 1 at some point during the construction of $\mathcal{M}(\mathcal{T})$. We wish to prove that for any $(s, \langle \Gamma, \mathcal{I}^L, \mathcal{I}^M \rangle) \in Y$ we have $\mathcal{M}(\mathcal{T}), s \models \Gamma$, where we write $\mathcal{M}(\mathcal{T}), s \models \Gamma$ to mean $\mathcal{M}(\mathcal{T}), s \models \varphi$ for all $\varphi \in \Gamma$. Note that if we can prove this, then it follows that $\mathcal{M}(\mathcal{T}), s_{\mathcal{T}} \models \varphi$ since $(s_{\mathcal{T}}, \langle \{\varphi\}, \emptyset, (0, 0) \rangle) \in Y$.

Let (s, m) be an arbitrary element of Y and let l be the length of the longest path from m to a leaf. We will prove, by induction on l , that $\mathcal{M}(\mathcal{T}), s \models \Gamma$ where $m = \langle \Gamma, \mathcal{I}^L, \mathcal{I}^M \rangle$.

$l = 0$: In this case, m is a leaf. Hence Γ only contains literals, and by construction we have $p \in \ell(s)$ if and only if $p \in \Gamma$. Since m is consistent, we thus get $\mathcal{M}(\mathcal{T}), s \models \Gamma$.

$l > 0$: In this case we consider the different rules that may be applied to m .

(\wedge) : We have

$$(\wedge) \frac{m = \langle \Gamma \cup \{\varphi_1 \wedge \varphi_2\}, \mathcal{I}^L, \mathcal{I}^M \rangle}{m' = \langle \Gamma \cup \{\varphi_1, \varphi_2\}, \mathcal{I}^L, \mathcal{I}^M \rangle}$$

By induction hypothesis we get $\mathcal{M}(\mathcal{T}), s \models \Gamma \cup \{\varphi_1, \varphi_2\}$. This implies that $\mathcal{M}(\mathcal{T}), s \models \varphi_1$ and $\mathcal{M}(\mathcal{T}), s \models \varphi_2$, so $\mathcal{M}(\mathcal{T}), s \models \Gamma \cup \{\varphi_1 \wedge \varphi_2\}$.

$(\neg\wedge)$: We have

$$(\neg\wedge) \frac{m = \langle \Gamma \cup \{\neg(\varphi_1 \wedge \varphi_2)\}, \mathcal{I}^L, \mathcal{I}^M \rangle}{m_1 = \langle \Gamma \cup \{\neg\varphi_1\}, \mathcal{I}^L, \mathcal{I}^M \rangle \quad m_2 = \langle \Gamma \cup \{\neg\varphi_2\}, \mathcal{I}^L, \mathcal{I}^M \rangle}$$

We have three cases to consider; either m_1 is included in \mathcal{T}' , m_2 is included in \mathcal{T}' , or both m_1 and m_2 are included in \mathcal{T}' . If m_1 is included in \mathcal{T}' we get, by the induction hypothesis, that $\mathcal{M}(\mathcal{T}), s \models \Gamma \cup \{\neg\varphi_1\}$ implying that $\mathcal{M}(\mathcal{T}), s \not\models \varphi_1$. If m_2 is included in \mathcal{T}' we get, by the induction hypothesis, that $\mathcal{M}(\mathcal{T}), s \models \Gamma \cup \{\neg\varphi_2\}$ implying that $\mathcal{M}(\mathcal{T}), s \not\models \varphi_2$. In either case we get that $\mathcal{M}(\mathcal{T}), s \not\models \varphi_1 \wedge \varphi_2$ and $\mathcal{M}(\mathcal{T}), s \models \Gamma$, and therefore $\mathcal{M}(\mathcal{T}), s \models \Gamma \cup \{\neg(\varphi_1 \wedge \varphi_2)\}$. The last case follows trivially from the preceding arguments.

$(\neg\neg)$: We have

$$(\neg\neg) \frac{m = \langle \Gamma \cup \{\neg\neg\varphi'\}, \mathcal{I}^L, \mathcal{I}^M \rangle}{m' = \langle \Gamma \cup \{\varphi'\}, \mathcal{I}^L, \mathcal{I}^M \rangle}$$

By induction hypothesis we know that $\mathcal{M}(\mathcal{T}), s \models \Gamma \cup \{\varphi'\}$, so $\mathcal{M}(\mathcal{T}), s \models \Gamma \cup \{\neg\neg\varphi'\}$.

(mod): We have

$$(\text{mod}) \frac{m = \langle \Gamma \cup \{N_{r_1}^1 \varphi_1, \dots, N_{r_n}^n \varphi_n\} \cup \{\neg O_{r'_1}^1 \varphi'_1, \dots, \neg O_{r'_n}^n \varphi'_n\}, \mathcal{I}^L, \mathcal{I}^M \rangle}{m_1 = \langle \{\psi_1\}, \mathcal{I}_1^L, \mathcal{I}_1^M \rangle \quad \dots \quad m_k = \langle \{\psi_k\}, \mathcal{I}_k^L, \mathcal{I}_k^M \rangle}$$

Γ must consist only of literals, because otherwise the (mod) rule could not be used. As in the case for $l = 0$, we then get $\mathcal{M}(\mathcal{T}), s \models \Gamma$ since m is consistent. Let $\Psi = \{\psi_1, \dots, \psi_k\}$, and for any $1 \leq j \leq k$, let $\mathcal{I}_j^L = \{a_j, b_j\}$ and $\mathcal{I}_j^M = \{c_j, d_j\}$. By the induction hypothesis, we know that $\mathcal{M}(\mathcal{T}), s_j \models \psi_j$ for all $j \in \{1, \dots, k\}$, and, by construction, s_j is the only successor of s that satisfies ψ_j . Now consider a formula $N_{r_i}^i \varphi_i$. There must exist a subset $\Psi_{\varphi_i} \subseteq \Psi$ such that $\theta(s)(\llbracket \varphi_i \rrbracket) = \theta(s)\left(\bigcup_{\psi' \in \Psi_{\varphi_i}} \llbracket \psi' \rrbracket\right)$.

We first consider the case where $N^i = L$. Because Ψ_{φ_i} is finite, there exists $\psi'_j \in \Psi_{\varphi_i}$ such that $\theta^-(s)(\llbracket \varphi_i \rrbracket) = \theta^-(s)(\llbracket \psi'_j \rrbracket)$, implying the existence of $\psi_j \in \Psi$ such that $\theta^-(s)(\llbracket \varphi_i \rrbracket) = \theta^-(s)(\llbracket \psi_j \rrbracket) = a_j$. We must have $a_j \geq r_i$ implying $\theta^-(s)(\llbracket \varphi_i \rrbracket) \geq r_i$, and thus $\mathcal{M}(\mathcal{T}), s \models L_{r_i} \varphi_i$. In the case where $N^i = M$ we can, similarly to the previous case, find $\psi_j \in \Psi$ such that $\theta^+(s)(\llbracket \varphi_i \rrbracket) = \theta^+(s)(\llbracket \psi_j \rrbracket)$, and we know that $d_i \neq \infty$ implying

$$\theta^+(s)(\llbracket \psi_j \rrbracket) = \max \left\{ a_j, \frac{d_j - c_j}{2} + c_j \right\} \leq d_j \leq r_i.$$

Therefore, $\theta^+(s)(\llbracket \varphi \rrbracket) \leq r_i$ and thus $\mathcal{M}(\mathcal{T}), s \models M_{r_i} \varphi_i$.

Lastly we consider a formula $\neg O_{r'_i}^i \varphi'_i$. If there is no $\psi_j \in \Psi$ such that $\models \psi_j \rightarrow \varphi'_i$, then, by the construction of $\mathcal{M}(\mathcal{T})$, there is no successor s' of s such that $\mathcal{M}(\mathcal{T}), s \models \varphi'_i$. Therefore, $\theta^-(s)(\llbracket \varphi'_i \rrbracket) = \infty$ and $\theta^+(s)(\llbracket \varphi'_i \rrbracket) = -\infty$, and thus $\mathcal{M}(\mathcal{T}), s \models \neg O_{r'_i}^i \varphi'_i$ is trivially satisfied for $O^i \in \{L, M\}$. Suppose $\models \psi'_j \rightarrow \varphi'_i$ for some $\psi'_j \in \Psi$. We first consider the case where $O^i = L$. There must exist $\psi_j \in \Psi$ such that $\theta^-(s)(\llbracket \varphi'_i \rrbracket) = \theta^-(s)(\llbracket \psi_j \rrbracket) = a_j$. By the assumption that \mathcal{T} is successful, we must have that m_j is consistent. Therefore, $a_j < b_j \leq r'_i$ implying $\theta^-(s)(\llbracket \varphi'_i \rrbracket) < r'_i$, and thus $\mathcal{M}(\mathcal{T}), s \models \neg L_{r'_i} \varphi'_i$. In the case where $O^i = M$ we must be able to find $\psi_j \in \Psi$ such that $\theta^+(s)(\llbracket \varphi'_i \rrbracket) = \theta^+(s)(\llbracket \psi_j \rrbracket)$. We have to consider $d_j = \infty$ and $d_j \neq \infty$ separately. If $d_j = \infty$ we have

$$\theta^+(s)(\llbracket \psi_j \rrbracket) = \max\{a_j, c_j + 1\} > c_j \geq r'_i.$$

If $d_j \neq \infty$ we have

$$\theta^+(s)(\llbracket \psi_j \rrbracket) = \max\left\{a_j, \frac{d_j - c_j}{2} + c_j\right\} > c_j \geq r'_i.$$

In either case we have that $\theta(s)(\llbracket \varphi'_i \rrbracket) > r'_i$ and therefore $\mathcal{M}(\mathcal{T}), s \models \neg M_{r'_i} \varphi'_i$. \square

Lemma 5.6. *Let \mathcal{T}_1 and \mathcal{T}_2 be tableaux for φ . Then it holds that \mathcal{T}_1 is successful if and only if \mathcal{T}_2 is successful.*

Proof. Assume that \mathcal{T}_1 is a successful tableau. Let \mathcal{T}'_1 be a subtree of \mathcal{T}_1 which witnesses the fact that \mathcal{T}_1 is successful. If \mathcal{T}'_1 is also a subtree of \mathcal{T}_2 , then we are done. If not, let d be the smallest number such that \mathcal{T}'_1 differs at depth d from any subtree of \mathcal{T}_2 with the same root as \mathcal{T}_2 . Note that we must have $d > 0$ because \mathcal{T}'_1 and \mathcal{T}_2 have the same root. Denote by $\mathcal{T}'_1|_n$ the restriction of \mathcal{T}'_1 to depth n . Then $\mathcal{T}'_1|_{d-1}$ is a subtree of \mathcal{T}_2 .

At this point we note that \mathcal{T}_1 and \mathcal{T}_2 contain the same terminal nodes. To see this, let the level k terminal nodes be those terminal nodes that can be reached from the root by going through $k - 1$ terminal nodes. We now argue that the level k terminal nodes of \mathcal{T}_1 and \mathcal{T}_2 are the same by induction on k .

$k = 1$: The level 1 terminal nodes of \mathcal{T}_1 and \mathcal{T}_2 must be the same, since they are all constructed by applying the (\wedge) , $(\neg\wedge)$, or $(\neg\neg)$ rules to the root node $\langle\{\varphi\}, [0, 0], [0, 0]\rangle$.

$k > 1$: Since the level $k - 1$ terminal nodes of \mathcal{T}_1 and \mathcal{T}_2 are the same, they must also have the same children, which are constructed from the (mod) rule. Hence each level k terminal node is constructed by applying the (\wedge) , $(\neg\wedge)$, or $(\neg\neg)$ rules to a child of one of the level $k - 1$ terminal nodes, so they are also the same in \mathcal{T}_1 and \mathcal{T}_2 .

Now let X be the set of all terminal nodes that are in \mathcal{T}'_1 at depth d or below. Since every node in X is a node in \mathcal{T}_1 , it must also be a node in \mathcal{T}_2 . Furthermore, every node in X is reachable in \mathcal{T}_2 from $\mathcal{T}'_1|_{d-1}$. Hence, if we extend $\mathcal{T}'_1|_{d-1}$ to include all paths in \mathcal{T}_2 leading from the leaves of $\mathcal{T}'_1|_{d-1}$ to an element in X , then this extension is a subtree of \mathcal{T}_2 . Denote this extension by \mathcal{T}'_2 .

Finally we argue that \mathcal{T}'_2 is a witness for the fact that \mathcal{T}_2 is successful by checking the three conditions of Definition 5.4. Every leaf of \mathcal{T}'_2 is also a leaf in \mathcal{T}_2 , since all the leaves of \mathcal{T}'_2 are elements of X . This takes care of the first condition. If m_i is a child of the modal node m in \mathcal{T}_2 , and m is included in \mathcal{T}'_2 , then m is also a modal node in \mathcal{T}'_1 , and hence m_i must be included in \mathcal{T}'_1 . This means that there is a terminal node m'_1 in \mathcal{T}'_1 which is reached by m_i . Hence, if m_i is not included in \mathcal{T}'_2 , then the terminal node m'_1 can not be reached in \mathcal{T}'_2 , but this contradicts how \mathcal{T}'_2 was constructed. Therefore m_i must also be included in \mathcal{T}'_2 , so the second condition is satisfied. The last condition is satisfied because every terminal

node in \mathcal{T}'_2 is also a terminal node in \mathcal{T}'_1 , and we know that every terminal node in \mathcal{T}'_1 is consistent. \square

Lemma 5.7. *φ is satisfiable if and only if there exists a successful tableau for φ .*

Proof. (\implies) Assume φ is satisfiable, meaning that $\mathcal{M}, s \models \varphi$ for some $\mathcal{M} = (S, \rightarrow, \ell)$ and $s \in S$.

Let \mathcal{T} be a tableau for φ , and note that such a tableau always exists by applying the tableau rules to $\langle \{\varphi\}, [0, 0], [0, 0] \rangle$. Now construct a marking $\mathfrak{M} \subseteq S \times \mathcal{T}$ as follows.

- $(s, r) \in \mathfrak{M}$ where r is the root of \mathcal{T} .
- If $(s', m) \in \mathfrak{M}$ and (\wedge) or $(\neg\neg)$ was applied to m , add (s', m') to \mathfrak{M} , where m' is the child of m .
- If $(s', m) \in \mathfrak{M}$ and $(\neg\wedge)$ was applied to m , meaning that

$$(\neg\wedge) \frac{m = \langle \Gamma \cup \{\neg(\varphi_1 \wedge \varphi_2)\}, \mathcal{I}^L, \mathcal{I}^M \rangle}{m_1 = \langle \Gamma \cup \{\neg\varphi_1\}, \mathcal{I}^L, \mathcal{I}^M \rangle \quad m_2 = \langle \Gamma \cup \{\neg\varphi_2\}, \mathcal{I}^L, \mathcal{I}^M \rangle}$$

then add (s', m_1) to \mathfrak{M} if $s' \in \llbracket \neg\varphi_1 \rrbracket$ and add (s', m_2) to \mathfrak{M} if $s' \in \llbracket \neg\varphi_2 \rrbracket$.

- If $(s', m) \in \mathfrak{M}$ and (mod) was applied to m , meaning that

$$(\text{mod}) \frac{m = \langle \Gamma \cup \{N_{r_1}^1 \varphi_1, \dots, N_{r_n}^n \varphi_n\} \cup \{\neg O_{r'_1}^1 \varphi'_1, \dots, \neg O_{r'_n}^n \varphi'_n\}, \mathcal{I}^L, \mathcal{I}^M \rangle}{m_1 = \langle \{\psi_1\}, \mathcal{I}_1^L, \mathcal{I}_1^M \rangle \quad \dots \quad m_k = \langle \{\psi_k\}, \mathcal{I}_k^L, \mathcal{I}_k^M \rangle}$$

then add (t', m_i) to \mathfrak{M} if $t' \in \llbracket \psi_i \rrbracket$ and $s' \xrightarrow{r} t'$ for some $r \in \mathbb{R}_{\geq 0}$.

We will first argue that for any $(s', \langle \Gamma, \mathcal{I}^L, \mathcal{I}^M \rangle) \in \mathfrak{M}$ we have $\mathcal{M}, s' \models \Gamma$, meaning $\mathcal{M}, s' \models \varphi'$ for all $\varphi' \in \Gamma$. We prove this by induction on the depth d of m .

$d = 0$: We have $(s', \langle \Gamma, \mathcal{I}^L, \mathcal{I}^M \rangle) = (s, r) = (s, \langle \{\varphi\}, [0, 0], [0, 0] \rangle)$, and by assumption we get $\mathcal{M}, s \models \varphi$.

$d > 0$: We consider which rule was applied to the parent of m .

(\wedge) :

$$(\wedge) \frac{m' = \langle \Gamma \cup \{\varphi_1 \wedge \varphi_2\}, \mathcal{I}^L, \mathcal{I}^M \rangle}{m = \langle \Gamma \cup \{\varphi_1, \varphi_2\}, \mathcal{I}^L, \mathcal{I}^M \rangle}$$

By induction hypothesis, we have $\mathcal{M}, s' \models \Gamma \cup \{\varphi_1 \wedge \varphi_2\}$, so $\mathcal{M}, s' \models \varphi_1$ and $\mathcal{M}, s' \models \varphi_2$, and hence $\mathcal{M}, s' \models \Gamma \cup \{\varphi_1, \varphi_2\}$.

$(\neg\wedge)$:

$$(\neg\wedge) \frac{m' = \langle \Gamma \cup \{\neg(\varphi_1 \wedge \varphi_2)\}, \mathcal{I}^L, \mathcal{I}^M \rangle}{m_1 = \langle \Gamma \cup \{\neg\varphi_1\}, \mathcal{I}^L, \mathcal{I}^M \rangle \quad m_2 = \langle \Gamma \cup \{\neg\varphi_2\}, \mathcal{I}^L, \mathcal{I}^M \rangle}$$

If $m = m_1$, then by the way \mathfrak{M} was constructed we get $\mathcal{M}, s' \models \neg\varphi_1$, and hence by induction hypothesis, $\mathcal{M}, s' \models \Gamma \cup \{\neg\varphi_1\}$. Likewise we get $\mathcal{M}, s' \models \Gamma \cup \{\neg\varphi_2\}$ if $m = m_2$.

$(\neg\neg)$:

$$(\neg\neg) \frac{m' = \langle \Gamma \cup \{\neg\neg\varphi'\}, \mathcal{I}^L, \mathcal{I}^M \rangle}{m = \langle \Gamma \cup \{\varphi'\}, \mathcal{I}^L, \mathcal{I}^M \rangle}$$

By induction hypothesis we have $\mathcal{M}, s' \models \Gamma \cup \{\neg\neg\varphi'\}$, which is equivalent to $\mathcal{M}, s' \models \Gamma \cup \{\varphi'\}$.

(mod):

$$\text{(mod)} \frac{m' = \langle \Gamma \cup \{N_{r_1}^1 \varphi_1, \dots, N_{r_n}^n \varphi_n\} \cup \{\neg O_{r'_1}^1 \varphi'_1, \dots, \neg O_{r'_n}^n \varphi'_{n'}\}, \mathcal{I}^L, \mathcal{I}^M \rangle}{m_1 = \langle \{\psi_1\}, \mathcal{I}_1^L, \mathcal{I}_1^M \rangle \quad \dots \quad m_k = \langle \{\psi_k\}, \mathcal{I}_k^L, \mathcal{I}_k^M \rangle}$$

We must have $m = m_i$ for some $1 \leq i \leq k$. By construction of \mathfrak{M} we know that $\mathcal{M}, m_i \models \psi_i$.

Now let \mathcal{T}' be the subtree of \mathcal{T} consisting of those nodes m where there exists a state s' such that $(s', m) \in \mathfrak{M}$. We will now prove that \mathcal{T}' satisfies the three conditions in Definition 5.4.

For the first condition we prove the contrapositive: If m is not a leaf in \mathcal{T} , then it is not a leaf in \mathcal{T}' . Hence we assume that m is not a leaf in \mathcal{T} . If m is not a node in \mathcal{T}' , then it is also not a leaf node in \mathcal{T}' . If m is a node in \mathcal{T}' , then there must exist some state s' such that $(s', m) \in \mathfrak{M}$. We now consider which rule was applied to m in \mathcal{T} .

(\wedge) or (\neg): In these cases, m has a child m' in \mathcal{T} , and by construction of \mathfrak{M} , we get $(s', m') \in \mathfrak{M}$, so m' is a child of m in \mathcal{T}' .

($\neg \wedge$):

$$\text{($\neg \wedge$)} \frac{m = \langle \Gamma \cup \{\neg(\varphi_1 \wedge \varphi_2)\}, \mathcal{I}^L, \mathcal{I}^M \rangle}{m_1 = \langle \Gamma \cup \{\neg\varphi_1\}, \mathcal{I}^L, \mathcal{I}^M \rangle \quad m_2 = \langle \Gamma \cup \{\neg\varphi_2\}, \mathcal{I}^L, \mathcal{I}^M \rangle}$$

We know that $\mathcal{M}, s' \models \Gamma \cup \{\neg(\varphi_1 \wedge \varphi_2)\}$, so we must have $\mathcal{M}, s' \models \neg\varphi_1$ or $\mathcal{M}, s' \models \neg\varphi_2$. By construction of \mathfrak{M} , this means that $(s', m_1) \in \mathfrak{M}$ or $(s', m_2) \in \mathfrak{M}$, and hence m_1 or m_2 must be a child of m in \mathcal{T}' .

(mod):

$$\text{(mod)} \frac{m = \langle \Gamma \cup \{N_{r_1}^1 \varphi_1, \dots, N_{r_n}^n \varphi_n\} \cup \{\neg O_{r'_1}^1 \varphi'_1, \dots, \neg O_{r'_n}^n \varphi'_{n'}\}, \mathcal{I}^L, \mathcal{I}^M \rangle}{m_1 = \langle \{\psi_1\}, \mathcal{I}_1^L, \mathcal{I}_1^M \rangle \quad \dots \quad m_k = \langle \{\psi_k\}, \mathcal{I}_k^L, \mathcal{I}_k^M \rangle}$$

For each m_i there must exist some j such that $N_{r_j}^j \varphi_j = N_{r_j}^j \psi_i$. Then we know that $\mathcal{M}, s' \models N_{r_j}^j \psi_i$, and hence $\theta^-(s')(\llbracket \psi_i \rrbracket) \geq r_j$ or $\theta^+(s')(\llbracket \psi_i \rrbracket) \leq r_j$. In either case there must exist some $t' \in \llbracket \psi_i \rrbracket$ such that $s' \xrightarrow{r} t'$ for some r . Hence $(t', m_i) \in \mathfrak{M}$ and m_i is a child of m in \mathcal{T}' .

For the second condition, let $(s', m) \in \mathfrak{M}$ where m is a modal node, meaning that

$$\text{(mod)} \frac{m = \langle \Gamma \cup \{N_{r_1}^1 \varphi_1, \dots, N_{r_n}^n \varphi_n\} \cup \{\neg O_{r'_1}^1 \varphi'_1, \dots, \neg O_{r'_n}^n \varphi'_{n'}\}, \mathcal{I}^L, \mathcal{I}^M \rangle}{m_1 = \langle \{\psi_1\}, \mathcal{I}_1^L, \mathcal{I}_1^M \rangle \quad \dots \quad m_k = \langle \{\psi_k\}, \mathcal{I}_k^L, \mathcal{I}_k^M \rangle}$$

For every ψ_i we must have $N_{r_j}^j \varphi_j = N_{r_j}^j \psi_i$ for some j , so $\mathcal{M}, s' \models N_{r_j}^j \psi_i$, which implies that there exists $t' \in \llbracket \psi_i \rrbracket$ such that $s' \xrightarrow{r} t'$ for some r . Hence we get $(t', m_i) \in \mathfrak{M}$. Since this holds for any i , we get that every m_i is included in \mathcal{T}' .

For the third condition, let $m = \langle \Gamma, \mathcal{I}^L, \mathcal{I}^M \rangle$ be a terminal node in \mathcal{T}' . We check the conditions of Definition 5.3. There must exist a state s' such that $(s', m) \in \mathfrak{M}$, which means that $\mathcal{M}, s' \models \Gamma$. Hence s' satisfies all the literals in Γ , which can only happen if the first condition is satisfied. For the second condition, note that $[0, 0]$ is a consistent interval, and every interval constructed by the (mod) rule is also consistent, so \mathcal{I}^L and \mathcal{I}^M must be consistent. Hence it remains to check the third condition. Assume that

$$\mathcal{I}^L = \lceil_1 a, b \rfloor_1 \quad \text{and} \quad \mathcal{I}^M = \lceil_2 c, d \rfloor_2.$$

Now, either $\mathcal{I}^L = \mathcal{I}^M = [0, 0]$, in which case clearly $a \leq b$, $\lceil_1 = [$, and $\rfloor_2 =]$, or there exists a modal node m' in \mathcal{T}' such that m can be reached from m' . Let

$$m^* = \langle \Gamma^* \cup \{N_{r_1}^1 \varphi_1, \dots, N_{r_n}^n \varphi_n\} \cup \{\neg O_{r'_1}^1 \varphi'_1, \dots, \neg O_{r'_n}^n \varphi'_n\}, \mathcal{I}_*^L, \mathcal{I}_*^M \rangle$$

be the modal node in \mathcal{T}' with greatest depth from which m can be reached. Then m^* must have a child $m_i^* = \langle \{\psi_i\}, \mathcal{I}_i^L, \mathcal{I}_i^M \rangle$ where

$$\mathcal{I}^L = \mathcal{I}_i^L \quad \text{and} \quad \mathcal{I}^M = \mathcal{I}_i^M.$$

If $\mathcal{I}_i^M = (\max \mathbb{M}_i^-, \infty)$ or $\mathcal{I}_i^M = [0, \infty)$, then clearly $a < d = \infty$. Otherwise, if $\mathcal{I}_i^L = [0, \min \mathbb{L}_i^-)$ or $\mathcal{I}_i^L = [0, \infty)$ and $\mathcal{I}_i^M = (\max \mathbb{M}_i^-, \min \mathbb{M}_i^+)$ or $\mathcal{I}_i^M = [0, \min \mathbb{M}_i^+)$, then $0 = a \leq d$, $\lceil_1 = [$, and $\rfloor_2 =]$. Otherwise, the only possibility left is that $\lceil_1 = [$, $a = \max \mathbb{L}_i^+$, $\rfloor_2 =]$, and $d = \min \mathbb{M}_i^+$. We must show that $a \leq d$. Assume towards a contradiction that $a = \max \mathbb{L}_i^+ > \min \mathbb{M}_i^+ = d$. Then, by the definition of \mathbb{L}_i^+ and \mathbb{M}_i^+ , there exist j_1 and j_2 such that

$$\begin{aligned} L_{r_1} \varphi_{j_1} &= N_{r_1}^{j_1} \varphi_{j_1} & \text{and} & \quad \varphi_{j_1} \in \uparrow_{\Gamma'}(\psi_i) \\ M_{r_2} \varphi_{j_2} &= N_{r_2}^{j_2} \varphi_{j_2} & \text{and} & \quad \varphi_{j_2} \in \uparrow_{\Gamma'}(\psi_i) \end{aligned}$$

with $r_1 > r_2$. Because m^* is a node in \mathcal{T}' , there must exist a state s^* such that $(s^*, m^*) \in \mathfrak{M}$, which implies that $\mathcal{M}, s^* \models L_{r_1} \varphi_{j_1}$ and $\mathcal{M}, s^* \models M_{r_2} \varphi_{j_2}$. This gives us

$$\theta^-(s^*) (\llbracket \psi_i \rrbracket) \geq \theta^-(s^*) (\llbracket \varphi_{j_1} \rrbracket) \geq r_1 > r_2 \geq \theta^+(s^*) (\llbracket \varphi_{j_2} \rrbracket) \geq \theta^+(s^*) (\llbracket \psi_i \rrbracket),$$

which is a contradiction. Hence $a \leq d$ and we are done.

(\Leftarrow) This follows from Lemma 5.5. □

Theorem 5.8. *The satisfiability problem for our logic is decidable.*

Proof. By Lemma 5.7, to decide whether a formula φ is satisfiable, it is enough to check whether there exists a successful tableau for φ . Furthermore, by Lemma 5.6 it is enough to only check a single tableau for φ : If the tableau is successful, then all tableaux for φ are successful, and if it is not successful, then no tableau for φ is successful.


One can construct such a tableau for φ by applying the tableau rules of Table 2 to the tuple $\langle \{\varphi\}, [0, 0], [0, 0] \rangle$ until no more rules can be applied. We will now argue that there is an effective procedure for constructing such a tableau by induction on the modal depth of φ .

$md(\varphi) = 0$: In this case, the (mod) rule is never used when constructing the tableau. Hence the procedure proceeds by syntactically checking which rules can be used at a given moment, and choosing a valid rule to apply.

$md(\varphi) > 0$: In this case we proceed as for the case where $md(\varphi) = 0$, except that now the (mod) rule may also be applied, in which case we need to be able to compute the ψ_i , \mathcal{I}_i^L and \mathcal{I}_i^M . The difficulty lies in computing the set $\{\psi_1, \dots, \psi_m\} = \min(\mathcal{L}(\Gamma'))$, where $\Gamma' = \{\varphi_1, \dots, \varphi_n\}$, and the sets

$$\begin{aligned} \mathbb{L}_i^+ &= \{r \mid L_r \varphi_j = N_{r_j}^j \varphi_j \text{ for some } j \text{ and } \varphi_j \in \uparrow_{\Gamma'}(\psi_i)\} \\ \mathbb{M}_i^+ &= \{r \mid M_r \varphi_j = N_{r_j}^j \varphi_j \text{ for some } j \text{ and } \varphi_j \in \uparrow_{\Gamma'}(\psi_i)\} \\ \mathbb{L}_i^- &= \{r \mid L_r \varphi'_j = O_{r_j}^j \varphi'_j \text{ for some } j \text{ and } \models \psi_i \rightarrow \varphi'_j\} \\ \mathbb{M}_i^- &= \{r \mid M_r \varphi'_j = O_{r_j}^j \varphi'_j \text{ for some } j \text{ and } \models \psi_i \rightarrow \varphi'_j\}. \end{aligned}$$

However, note that all φ_i and φ'_i have modal depth less than $md(\varphi)$. Therefore, by induction hypothesis, we have an effective procedure to decide whether $\models \varphi_i \rightarrow \varphi_j$ and

FIGURE 6. The model $\mathcal{M}(\mathcal{T})$ for the successful tableau \mathcal{T} in Example 5.9.

$\models \varphi_i \leftrightarrow \varphi_j$, which is exactly what we need to compute the aforementioned sets. Given this we can compute the values needed for the intervals \mathcal{I}_i^L and \mathcal{I}_i^M .

The procedure for constructing a tableau for φ uses recursion on the modal depth of φ , $md(\varphi) = k$, in order to compute the sets $\{\psi_1, \dots, \psi_m\}$, \mathbb{L}_i^+ , \mathbb{M}_i^+ , \mathbb{L}_i^- , and \mathbb{M}_i^- . To compute these sets we must instantiate the procedure for constructing tableaux for formulae of modal depth $k - 1$, which again must instantiate the procedure for constructing tableaux for formulae of modal depth $k - 2$, and so on. The recursion stops when only the procedure for generating tableaux for formulae with modal depth zero is needed to construct the sets $\{\psi_1, \dots, \psi_m\}$, \mathbb{L}_i^+ , \mathbb{M}_i^+ , \mathbb{L}_i^- , and \mathbb{M}_i^- .

Thus, for any $k \in \mathbb{N}$, there exists a procedure for generating a tableau for any formula φ with $md(\varphi) = k$. Because all formulae are finite they must have finite modal depth. Therefore, for any formula φ , there exists a procedure that generates a tableau for φ . \square

Example 5.9. Consider the formula $\varphi = \neg(\neg(L_2p_1 \wedge M_5L_1p_1) \wedge \neg M_2p_2)$. Using the tableau rules, we get the following tableau \mathcal{T} for φ .

$$\begin{array}{c}
(\neg\wedge) \frac{\langle \{\neg(\neg(L_2p_1 \wedge M_5L_1p_1) \wedge M_2p_2)\}, [0, 0], [0, 0] \rangle}{\langle \{\neg\neg(L_2p_1 \wedge M_5L_1p_1)\}, [0, 0], [0, 0] \rangle} \\
(\neg\neg) \frac{\langle \{\neg\neg(L_2p_1 \wedge M_5L_1p_1)\}, [0, 0], [0, 0] \rangle}{\langle \{L_2p_1 \wedge M_5L_1p_1\}, [0, 0], [0, 0] \rangle} \quad (\neg\neg) \frac{\langle \{\neg\neg M_2p_2\}, [0, 0], [0, 0] \rangle}{\langle \{M_2p_2\}, [0, 0], [0, 0] \rangle} \\
(\wedge) \frac{\langle \{L_2p_1 \wedge M_5L_1p_1\}, [0, 0], [0, 0] \rangle}{\langle \{L_2p_1, M_5L_1p_1\}, [0, 0], [0, 0] \rangle} \quad (\text{mod}) \frac{\langle \{M_2p_2\}, [0, 0], [0, 0] \rangle}{\langle \{p_2\}, [0, \infty), [0, 2] \rangle} \\
(\text{mod}) \frac{\langle \{L_2p_1, M_5L_1p_1\}, [0, 0], [0, 0] \rangle}{\langle \{p_1, L_1p_1\}, [2, \infty), [5, \infty) \rangle} \\
(\text{mod}) \frac{\langle \{p_1, L_1p_1\}, [2, \infty), [5, \infty) \rangle}{\langle \{p_1\}, [1, \infty), [0, \infty) \rangle}
\end{array}$$

In this case the tableau is successful, since all terminal nodes are consistent. In fact, there are three distinct subtrees witnessing this fact: one that chooses the left branch, one that chooses the right branch, and one that chooses both branches. In Figure 6 we show the resulting model $\mathcal{M}(\mathcal{T})$ for the witness that chooses the left branch.

Example 5.10. Consider the formula $\varphi = p_1 \wedge L_4p_1 \wedge \neg L_3p_1 \wedge L_2p_2$. Using the tableau rules, we get the following tableau \mathcal{T} for φ .

$$\begin{array}{c}
(\wedge) \frac{\langle \{p_1 \wedge L_4p_1 \wedge \neg L_3p_1 \wedge L_2p_2\}, [0, 0], [0, 0] \rangle}{\langle \{p_1, L_4p_1 \wedge \neg L_3p_1 \wedge L_2p_2\}, [0, 0], [0, 0] \rangle} \\
(\wedge) \frac{\langle \{p_1, L_4p_1 \wedge \neg L_3p_1 \wedge L_2p_2\}, [0, 0], [0, 0] \rangle}{\langle \{p_1, L_4p_1, \neg L_3p_1 \wedge L_2p_2\}, [0, 0], [0, 0] \rangle} \\
(\wedge) \frac{\langle \{p_1, L_4p_1, \neg L_3p_1 \wedge L_2p_2\}, [0, 0], [0, 0] \rangle}{\langle \{p_1, L_4p_1, \neg L_3p_1, L_2p_2\}, [0, 0], [0, 0] \rangle} \\
(\text{mod}) \frac{\langle \{p_1, L_4p_1, \neg L_3p_1, L_2p_2\}, [0, 0], [0, 0] \rangle}{\langle \{p_1\}, [4, 3), [0, \infty) \rangle \quad \langle \{p_2\}, [2, \infty), [0, \infty) \rangle}
\end{array}$$

In this case the interval $[4, 3)$ is not consistent, and hence the tableau is not successful, so we can conclude that φ is not satisfiable.

6. CONCLUDING REMARKS

Our contributions in this paper have been to define a new bisimulation relation for weighted transition systems, which relates those states that have similar behavior with respect to their minimum and maximum weights on transitions, as well as an accompanying modal logic to reason about the upper and lower bounds of weights on transitions. We have shown that this logic characterizes exactly those states that are bisimilar for image-finite systems. Furthermore, we have provided a complete axiomatization of our logic, and we have shown that it enjoys the finite model property. Lastly we have developed an algorithm based on the tableau method which decides the satisfiability of a formula in our logic and constructs a finite model for the formula if it is satisfiable.

This work could be extended in different ways. Since our logic is non-compact, strong completeness does not follow directly from weak completeness, and hence it would be interesting to explore a strong-complete axiomatization of the proposed logic. Such an axiomatization would need additional, infinitary axioms. Examples of such axioms would be $\{L_q\varphi \mid q < r\} \vdash L_r\varphi$ and $\{M_q\varphi \mid q < r\} \vdash M_r\varphi$, which are easily proven sound and describe the Archimedean property discussed in Theorem 4.11.

Although we have shown that our logic is expressive enough to capture bisimulation, it would also be of interest to extend our logic with a kind of fixed-point operator or standard temporal logic operators such as until in order to increase its expressivity, and hence its practical use. We envisage two ways in which such a logic could be given semantics: either by accumulating weights or by taking the maximum or minimum of weights. In the accumulating case in particular, one could also allow negative weights to model that the system gains resources.

Acknowledgements. We wish to thank the anonymous reviewers for their careful reading of our paper and for their invaluable comments that helped improve the paper. We are also grateful to Bingtian Xue for helpful discussions. This research was partially supported by the Danish FNU project 4181-00360, the ERC Advanced Grant LASSO: “Learning, Analysis, Synthesis and Optimization of Cyber Physical Systems” as well as the Sino-Danish Basic Research Center IDEA4CPS.

REFERENCES

- [ACD93] Rajeev Alur, Costas Courcoubetis, and David L. Dill. Model-checking in dense real-time. *Inf. Comput.*, 104(1):2–34, 1993.
- [BDP16] Parvaneh Babari, Manfred Droste, and Vitaly Perevoshchikov. Weighted register automata and weighted logic on data words. In Augusto Sampaio and Farn Wang, editors, *Theoretical Aspects of Computing - ICTAC 2016 - 13th International Colloquium, Taipei, Taiwan, ROC, October 24-31, 2016, Proceedings*, volume 9965 of *Lecture Notes in Computer Science*, pages 370–384, 2016.
- [BG09] Benedikt Bollig and Paul Gastin. Weighted versus probabilistic logics. In Volker Diekert and Dirk Nowotka, editors, *Developments in Language Theory, 13th International Conference, DLT 2009, Stuttgart, Germany, June 30 - July 3, 2009. Proceedings*, volume 5583 of *Lecture Notes in Computer Science*, pages 18–38. Springer, 2009.
- [BvBW06] P. Blackburn, J. F. A. K. van Benthem, and F. Wolter. *Handbook of Modal Logic*. Studies in Logic and Practical Reasoning. Elsevier Science, 2006.
- [CK16] Soumyodip Chakraborty and Joost-Pieter Katoen. On the satisfiability of some simple probabilistic logics. In Martin Grohe, Eric Koskinen, and Natarajan Shankar, editors, *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science, LICS '16, New York, NY, USA, July 5-8, 2016*, pages 56–65. ACM, 2016.

- [CLM11a] Luca Cardelli, Kim G. Larsen, and Radu Mardare. Continuous Markovian logic - from complete axiomatization to the metric space of formulas. In Marc Bezem, editor, *Computer Science Logic, 25th International Workshop / 20th Annual Conference of the EACSL, CSL 2011, September 12-15, 2011, Bergen, Norway, Proceedings*, volume 12 of *LIPICs*, pages 144–158. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2011.
- [CLM11b] Luca Cardelli, Kim G. Larsen, and Radu Mardare. Modular Markovian logic. In Luca Aceto, Monika Henzinger, and Jirí Sgall, editors, *Automata, Languages and Programming - 38th International Colloquium, ICALP 2011, Zurich, Switzerland, July 4-8, 2011, Proceedings, Part II*, volume 6756 of *Lecture Notes in Computer Science*, pages 380–391. Springer, 2011.
- [DG05] Manfred Droste and Paul Gastin. Weighted automata and weighted logics. In Luís Caires, Giuseppe F. Italiano, Luís Monteiro, Catuscia Palamidessi, and Moti Yung, editors, *Automata, Languages and Programming, 32nd International Colloquium, ICALP 2005, Lisbon, Portugal, July 11-15, 2005, Proceedings*, volume 3580 of *Lecture Notes in Computer Science*, pages 513–525. Springer, 2005.
- [DR06] Manfred Droste and George Rahonis. Weighted automata and weighted logics on infinite words. In Oscar H. Ibarra and Zhe Dang, editors, *Developments in Language Theory, 10th International Conference, DLT 2006, Santa Barbara, CA, USA, June 26-29, 2006, Proceedings*, volume 4036 of *Lecture Notes in Computer Science*, pages 49–58. Springer, 2006.
- [DV06] Manfred Droste and Heiko Vogler. Weighted tree automata and weighted logics. *Theor. Comput. Sci.*, 366(3):228–247, 2006.
- [Ési14] Zoltán Ésik. Axiomatizing weighted synchronization trees and weighted bisimilarity. *Theor. Comput. Sci.*, 534:2–23, 2014.
- [FH94] Ronald Fagin and Joseph Y. Halpern. Reasoning about knowledge and probability. *J. ACM*, 41(2):340–367, 1994.
- [Fic11] Ina Fichtner. Weighted picture automata and weighted logics. *Theory Comput. Syst.*, 48(1):48–78, 2011.
- [GH09] Steven Givant and Paul Halmos. *Introduction to Boolean Algebras*. Undergraduate Texts in Mathematics. Springer, 2009.
- [HLM⁺16] Mikkel Hansen, Kim Guldstrand Larsen, Radu Mardare, Mathias Ruggaard Pedersen, and Bingtian Xue. A complete approximation theory for weighted transition systems. In Martin Fränzle, Deepak Kapur, and Naijun Zhan, editors, *Dependable Software Engineering: Theories, Tools, and Applications - Second International Symposium, SETTA 2016, Beijing, China, November 9-11, 2016, Proceedings*, volume 9984 of *Lecture Notes in Computer Science*, pages 213–228, 2016.
- [HM01] Aviad Heifetz and Philippe Mongin. Probability logic for type spaces. *Games and Economic Behavior*, 35(1-2):31–53, 2001.
- [JL91] Bengt Jonsson and Kim Guldstrand Larsen. Specification and refinement of probabilistic processes. In *Proceedings of the Sixth Annual Symposium on Logic in Computer Science (LICS '91), Amsterdam, The Netherlands, July 15-18, 1991*, pages 266–277. IEEE Computer Society, 1991.
- [JLMX14] Samy Jaziri, Kim Guldstrand Larsen, Radu Mardare, and Bingtian Xue. Adequacy and complete axiomatization for timed modal logic. *Electr. Notes Theor. Comput. Sci.*, 308:183–210, 2014.
- [JLS12] Line Juhl, Kim G. Larsen, and Jirí Srba. Modal transition systems with weight intervals. *J. Log. Algebr. Program.*, 81(4):408–421, 2012.
- [KLMP13] Dexter Kozen, Kim G. Larsen, Radu Mardare, and Prakash Panangaden. Stone duality for Markov processes. In *28th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2013, New Orleans, LA, USA, June 25-28, 2013*, pages 321–330. IEEE Computer Society, 2013.
- [KMP13] Dexter Kozen, Radu Mardare, and Prakash Panangaden. Strong completeness for Markovian logics. In Krishnendu Chatterjee and Jirí Sgall, editors, *Mathematical Foundations of Computer Science 2013 - 38th International Symposium, MFCS 2013, Klosterneuburg, Austria, August 26-30, 2013. Proceedings*, volume 8087 of *Lecture Notes in Computer Science*, pages 655–666. Springer, 2013.
- [LM14] Kim G. Larsen and Radu Mardare. Complete proof systems for weighted modal logic. *Theor. Comput. Sci.*, 546:164–175, 2014.
- [LMX14a] Kim Guldstrand Larsen, Radu Mardare, and Bingtian Xue. Decidability and expressiveness of recursive weighted logic. In Andrei Voronkov and Irina Virbitskaite, editors, *Perspectives of System Informatics - 9th International Ershov Informatics Conference, PSI 2014, St. Petersburg,*

- Russia, June 24-27, 2014. Revised Selected Papers*, volume 8974 of *Lecture Notes in Computer Science*, pages 216–231. Springer, 2014.
- [LMX14b] Kim Guldstrand Larsen, Radu Mardare, and Bingtian Xue. A decidable recursive logic for weighted transition systems. In Gabriel Ciobanu and Dominique Méry, editors, *Theoretical Aspects of Computing - ICTAC 2014 - 11th International Colloquium, Bucharest, Romania, September 17-19, 2014. Proceedings*, volume 8687 of *Lecture Notes in Computer Science*, pages 460–476. Springer, 2014.
- [LMX15a] Kim G. Larsen, Radu Mardare, and Bingtian Xue. Alternation-free weighted mu-calculus: Decidability and completeness. *Electr. Notes Theor. Comput. Sci.*, 319:289–313, 2015.
- [LMX15b] Kim G. Larsen, Radu Mardare, and Bingtian Xue. Concurrent weighted logic. *J. Log. Algebr. Meth. Program.*, 84(6):884–897, 2015.
- [LMX18] Kim G. Larsen, Radu Mardare, and Bingtian Xue. On decidability of recursive weighted logics. *Soft Comput.*, 22(4):1085–1102, 2018.
- [LS91] Kim Guldstrand Larsen and Arne Skou. Bisimulation through probabilistic testing. *Inf. Comput.*, 94(1):1–28, 1991.
- [MCL12] Radu Mardare, Luca Cardelli, and Kim G. Larsen. Continuous Markovian logics - axiomatization and quantified metatheory. *Logical Methods in Computer Science*, 8(4), 2012.
- [Mei06] Ingmar Meinecke. Weighted logics for traces. In Dima Grigoriev, John Harrison, and Edward A. Hirsch, editors, *Computer Science - Theory and Applications, First International Computer Science Symposium in Russia, CSR 2006, St. Petersburg, Russia, June 8-12, 2006, Proceedings*, volume 3967 of *Lecture Notes in Computer Science*, pages 235–246. Springer, 2006.
- [Zho09] Chunlai Zhou. A complete deductive system for probability logic. *J. Log. Comput.*, 19(6):1427–1454, 2009.