Notions of Bisimulation for Heyting-Valued Modal Languages

Pantelis E. Eleftheriou[1]*, Costas D. Koutras[2], and Christos Nomikos[3]**

[1] Institut de Matemàtica, Universitat de Barcelona,
Gran Via de les Corts Catalanes 585, 08007 Barcelona, Spain
pelefthe@gmail.com

[2] Department of Computer Science and Technology, University of Peloponnese,
end of Karaiskaki Street, 22100 Tripolis, Greece
ckoutras@uop.gr

[3] Department of Computer Science, University of Ioannina,
45110 Ioannina, Greece
cnomikos@cs.uoi.gr

Abstract. We define notions of bisimulation for the family of Heyting-valued modal logics introduced by M. Fitting. In this family of logics, each modal language is built on an underlying space of truth values, a Heyting algebra $\mathcal{H}$. All the truth values are directly represented in the language, which is interpreted on relational frames with an $\mathcal{H}$-valued accessibility relation. We investigate the correct notion of bisimulation in this context: we define two variants of bisimulation relations and derive relative (to a truth value) modal equivalence results for bisimilar states. We further investigate game semantics for our bisimulation, Hennessy-Milner classes and other relevant properties. If the underlying algebra $\mathcal{H}$ is finite, Heyting-valued modal models can be equivalently reformulated to a form relevant to epistemic situations with many interrelated experts. Our definitions and results draw from this formulation, which is of independent interest to Knowledge Representation applications.

Key words: Modal Logic, Many-Valued Logic, Bisimulations.

1 Introduction

Bisimulation is a very rich concept which plays an important role in many areas of Computer Science, Logic and Set Theory. Its origins can be found in the analysis of Modal Logic but it was independently rediscovered by computer scientists in their efforts to understand concurrency. In Modal Logic, bisimulations were introduced by Johan van Benthem, under the name of p-relations or zigzag relations, in the course of his work on the correspondence theory of Modal Logic [vB83,vB84]. In Computer Science, bisimulations were introduced by Park

* Research supported by European Commission's Research Training Network MRTN-CT-2004-512234 (MODNET)

** Corresponding author

in [Par91] and Henessy and Milner in [HM85], in the course of investigating the notion of equivalence among processes (see [San07] for a tutorial on the history of bisimulations). In this context, bisimulations represent a fundamental notion of identity between process states and every language designed to capture the essential properties of processes should be blind for bisimilar states.

On the other hand, from the Modal Logic viewpoint, bisimulation is the correct notion of similarity between two modal models: modal formulas are unable to distinguish bisimilar points of the two models. But its importance lies far beyond; the celebrated van Benthem characterization theorem, published in the mid-'70s, states that invariance for bisimulation captures the essential property of the 'modal fragment' of first-order logic: *a first-order formula is invariant for bisimulation iff it is (equivalent to the) the syntactical translation of a modal formula* (see [BdRV01] for a nice exposition of this result and its consequences). The characterization theorem has generated an important stream of research in the analysis of logical languages. In particular, the bisimulation-based analysis of modal languages has been extensively studied in the Amsterdam school of modal logic [dR93,Ger99,MV03], and it has even been suggested that this notion is as important for modal logic as the notion of partial isomorphism has been for the model theory of classical logic (see the PhD thesis of M. de Rijke [dR93]). It is worth mentioning that bisimulations have been used beyond the realm of classical modal logic: see for instance the variant used to analyze since-until temporal languages [KdR97], M. Otto's work related to Finite Model Theory [Ott99] and J. Gerbrandy's dissertation [Ger99]. Bisimulations have been also used as a fundamental tool in the area of non-well founded set theory ([Acz88], see [BdRV01] for a few details and further references).

In this note, we address the question of what constitutes a suitable notion of bisimulation for the family of many-valued modal languages introduced by M. Fitting in the early '90s [Fit92,Fit91]. Each language of this family is built on an underlying space of truth values, a Heyting algebra $\mathcal{H}$. There exist three features that give these logics their distinctive character. The first one is syntactic: the elements of $\mathcal{H}$ are directly encoded in the language as special constants and this permits the formation of 'weak', uncertainty-oriented versions of the classical modal epistemic actions [Kou03,KNP02,KP02]. The second is semantic: the languages we discuss are interpreted on $\mathcal{H}$-labelled directed graphs which provide us a form of many-valued accessibility relation. Finally, the third one concerns the potential applicability of these logics in epistemic situations with multiple intelligent agents. More specifically, assuming that $\mathcal{H}$ is a finite Heyting algebra, these logics can be formulated in a way that expresses the epistemic consensus of many experts, interrelated through a binary 'dominance' relation [Fit92]. It is worth mentioning that, model-theoretically, every complete Heyting algebra can serve as the space of truth values. However, apart from the equivalent multiple-expert formulation of the logics, the finiteness assumption for $\mathcal{H}$ is essential for the elegant canonical model construction of [Fit92] which leads to a completeness theorem; note that this finiteness restriction seems to be also necessary for obtaining a many-valued analog of the ultrafilter extension construction [EK05].

In Section 2, we briefly expose the mathematical background of Heyting algebras, as well as the syntax and semantics of many-valued modal languages.

Section 3 contains the main definitions and results of the paper. We provide two notions of bisimulation and derive modal equivalence results. In Section 3.1, a rather strong notion of a *t-bisimulation* is defined for each $t \in \mathcal{H}$. This allows us to formulate simple, intuitive, Ehrenfeucht-Fraissé type bisimulation games through which one can easily define bounded t-bisimulations, as in the classical case. Also, an appropriate notion of *t-unravelling* is given, through which one gets a form of the celebrated *tree-model property*, considered to be critical for the 'robust decidability' of modal logics [Var97]. In Section 3.2, a rather involved notion of a *weak bisimulation* is defined, assuming our Heyting algebra $\mathcal{H}$ satisfies Property (1), given in the beginning of that section. This allows us to obtain an interesting notion of a Hennessy-Milner class of Heyting-valued modal models.

Both notions of bisimulations draw inspiration from the equivalent multiple-expert formulation of these logics, which is actually a mixture of Kripke modal and Kripke intuitionistic semantics. In Section 4, we review this alternative formulation and interpret our results in that context.

2 Many-Valued Modal Languages

In this section we provide the syntax and semantics of many-valued modal languages, as introduced in [Fit92], with only minor changes in the notation. To construct a modal language of this family, we first fix a Heyting algebra $\mathcal{H}$ which will serve as the space of truth values. Thus, we first briefly expose the basic definitions and properties of Heyting algebras, fixing also notation and terminology. We assume that the reader already has some familiarity with the elements of lattice theory and universal algebra. For more details the reader is referred to the classical texts [RS70,BD74].

Heyting Algebras A *lattice* $\mathcal{L}$ is a pair $\langle L, \leq \rangle$ consisting of a non-empty set L equipped with a partial-order relation $\leq$, such that every two-element subset $\{a, b\}$ of L has a *least upper bound* or *join*, denoted by $a \vee b$, and a *greatest lower bound* or *meet*, denoted by $a \wedge b$. A lattice $\mathcal{L}$ is *complete* if a join and a meet exist for *every* subset of $\mathcal{L}$. A *least* (or *bottom*) element of a lattice is denoted by $\bot$ and a *greatest* (or *top*) one by $\top$. An element $x \in L$ is *join-irreducible* if $x \neq \bot$ (in case L has a bottom element) and $x = a \vee b$ implies $x = a$ or $x = b$. We frequently use indexed sets and denote (possibly infinite) meets and joins by $\bigwedge_{t \in T} a_t$ and $\bigvee_{t \in T} a_t$. Some fairly obvious properties of infinite joins and meets, such as $\bigwedge_{t \in T} (a \wedge a_t) = a \wedge \bigwedge_{t \in T} a_t$ will be used, generally without comment.

A lattice $\mathcal{H} = \langle H, \leq \rangle$ with the additional property that, for every pair of elements $\langle a, b \rangle$, the set $\{x \mid a \wedge x \leq b\}$ has a greatest element, is called a *relatively pseudo-complemented lattice*. This element is denoted by $a \Rightarrow b$ and is called the *pseudo-complement of a relative to b*. A relatively pseudo-complemented lattice is always a topped ordered set [RS70]. It is not always the case that a relatively pseudo-complemented lattice has a least element. A relatively pseudo-complemented lattice $\mathcal{H}$ with a least element is called a **Heyting algebra** (HA) or a **pseudo-Boolean algebra**. It is known that the class of HAs includes the class of Boolean algebras and is included in the class of distributive lattices; both inclusions are proper. For finite lattices, the second inclusion becomes an equality: the class of finite HAs coincides with the class of finite distributive lattices [RS70]. The following lemma gathers some useful properties of relatively pseudocomplemented lattices that will be used in Section 3; whenever a possibly infinite join or meet is involved, it is assumed that it exists. The proof of its items can be found in [RS70,BD74]. Note also that the first item of the lemma can be equivalently considered as a definition of relative pseudo-complementation.

Lemma 1. *1.* $x \leq (a \Rightarrow b)$ *iff* $(x \wedge a) \leq b$
2. $(a \Rightarrow b) = \top$ *iff* $a \leq b$
3. *If* $a_1 \leq a_2$ *then* $(a_2 \Rightarrow b) \leq (a_1 \Rightarrow b)$
4. $c \wedge (a \Rightarrow b) = c \wedge \big((c \wedge a) \Rightarrow (c \wedge b)\big)$
5. $\bigvee_{t \in T} (a \wedge b_t) = a \wedge \bigvee_{t \in T} b_t$
6. $(a \vee b) \Rightarrow c = (a \Rightarrow c) \wedge (b \Rightarrow c)$

Syntax of Many-Valued Modal Languages Having fixed a complete Heyting algebra $\mathcal{H}$ we proceed to define the syntax of the modal language. The elements of $\mathcal{H}$ are directly represented in the language by special constants, called *propositional constants*, and we reserve lowercase letters (along with $\bot, \top$) to denote them. To facilitate notation, we use the same letter for the element of $\mathcal{H}$ and the constant which represents it in the language; context will clarify what is meant. Assuming also a set Φ of *propositional variables* (also called propositional letters) we define the many-valued modal language $L^{\mathcal{H}}_{\Box\Diamond}(\Phi)$ with the following BNF specification, where t ranges over elements of $\mathcal{H}$, P ranges over elements of Φ and A is a formula of $L^{\mathcal{H}}_{\Box\Diamond}(\Phi)$.

$$A \ ::= \ t \ \mid \ P \ \mid \ A_1 \vee A_2 \ \mid \ A_1 \wedge A_2 \ \mid \ A_1 \supset A_2 \ \mid \ \Box A \ \mid \ \Diamond A$$

As we will see below, the modal logics defined are in general bimodal, thus we need both modal operators. Note also that $\vee$ and $\wedge$ serve both as logical connectives, as well as lattice operation symbols but it should be clear by context what is meant. In the rest of the paper, we shall often omit Φ when possible and speak of the language $L^{\mathcal{H}}_{\Box\Diamond}$. A (non-classical) negation $\neg X$ can be defined as $(X \supset \bot)$.

Semantics of Many-Valued Modal Languages $L^{\mathcal{H}}_{\Box\Diamond}$ is interpreted on an interesting variant of a relational frame, which possesses a kind of Heyting-valued accessibility relation. Note that there have been other approaches in the literature for defining many-valued modal logics, but all of them have kept the essence of classical relational semantics intact (see [Fit92] for references). Given $L^{\mathcal{H}}_{\Box\Diamond}(\Phi)$, we define $\mathcal{H}$-modal frames and $\mathcal{H}$-modal models as follows:

Definition 1. An $\mathcal{H}$-*modal frame* for $L^{\mathcal{H}}_{\Box\Diamond}(\Phi)$ is a pair $\mathfrak{F} = \langle \mathfrak{S}, \mathfrak{g} \rangle$, where $\mathfrak{S}$ is a non-empty set of *states* and $\mathfrak{g} : \mathfrak{S} \times \mathfrak{S} \to \mathcal{H}$ is a total function mapping pairs of states to elements of $\mathcal{H}$.
An $\mathcal{H}$-*modal model* $\mathfrak{M} = \langle \mathfrak{S}, \mathfrak{g}, v \rangle$ is built on $\mathfrak{F}$ by providing a *valuation* v, that is a function $v : \mathfrak{S} \times (\mathcal{H} \cup \Phi) \to \mathcal{H}$ which assigns a $\mathcal{H}$-truth value to atomic formulae in each state, such that $v(\mathfrak{s}, t) = t$, for every $\mathfrak{s} \in \mathfrak{S}$ and $t \in \mathcal{H}$. In other words, the propositional constants are always mapped to 'themselves'.

In the sequel, we shall often omit the adjective 'modal' and talk simply of $\mathcal{H}$-frames and $\mathcal{H}$-models.
The valuation v extends to all the formulae of $L^{\mathcal{H}}_{\Box\Diamond}(\Phi)$ in a standard recursive fashion:

Definition 2. Let $\mathfrak{M} = \langle \mathfrak{S}, \mathfrak{g}, v \rangle$ be an $\mathcal{H}$-*model* and $\mathfrak{s}$ a state of $\mathfrak{S}$. The extension of the valuation v to the whole language $L^{\mathcal{H}}_{\Box\Diamond}(\Phi)$ is given by the following clauses;

- $v(\mathfrak{s}, A \wedge B) = v(\mathfrak{s}, A) \wedge v(\mathfrak{s}, B)$
- $v(\mathfrak{s}, A \vee B) = v(\mathfrak{s}, A) \vee v(\mathfrak{s}, B)$
- $v(\mathfrak{s}, A \supset B) = v(\mathfrak{s}, A) \Rightarrow v(\mathfrak{s}, B)$
- $v(\mathfrak{s}, \Box A) = \bigwedge_{\mathfrak{t} \in \mathfrak{S}} \big(\mathfrak{g}(\mathfrak{s}, \mathfrak{t}) \Rightarrow v(\mathfrak{t}, A)\big)$
- $v(\mathfrak{s}, \Diamond A) = \bigvee_{\mathfrak{t} \in \mathfrak{S}} \big(\mathfrak{g}(\mathfrak{s}, \mathfrak{t}) \wedge v(\mathfrak{t}, A)\big)$

The operators of *necessity* ($\Box$) and *possibility* ($\Diamond$) are not each other's dual, unless $\mathcal{H}$ is a Boolean algebra [Fit92]. Note also that all the definitions above collapse to the familiar ones from the classical case, in the case of the classical language $L^{\mathbf{2}}_{\Box\Diamond}$, where $\mathbf{2}$ is the lattice of two-valued classical logic.

3 Bisimulations for Many-Valued Modal Languages

In this section, we define two suitable general notions of bisimulation for a language $L^{\mathcal{H}}_{\Box\Diamond}$ of the family defined in the previous section. Before proceeding, we have to define a refined notion of modal truth invariance which fits our aims and which also has an interesting interpretation in the multiple-expert context. Note that the following notion is trivial for $t = \bot$.

Definition 3 (t-invariance). Let $\mathfrak{M} = \langle \mathfrak{S}, \mathfrak{g}, v \rangle$ and $\mathfrak{M}' = \langle \mathfrak{S}', \mathfrak{g}', v' \rangle$ be $\mathcal{H}$-models for $L^{\mathcal{H}}_{\Box\Diamond}(\Phi)$, $\mathfrak{s} \in \mathfrak{S}$ and $\mathfrak{s}' \in \mathfrak{S}'$ two states and $t \in \mathcal{H}$ a truth value ($t \neq \bot$). We say that modal truth is *t-invariant for the transition between $\mathfrak{s}$ and $\mathfrak{s}'$* if for every $X \in L^{\mathcal{H}}_{\Box\Diamond}(\Phi)$

$$t \wedge v(\mathfrak{s}, X) = t \wedge v'(\mathfrak{s}', X)$$

3.1 Strong Bisimulation for Many-Valued Modal Languages

The following definition captures the idea of moving 'back and forth' between two $\mathcal{H}$-models by matching steps ('modulo' t) in both directions.

Definition 4 (t-bisimulations). Given two $\mathcal{H}$-models $\mathfrak{M} = \langle \mathfrak{S}, \mathfrak{g}, v \rangle$ and $\mathfrak{M}' = \langle \mathfrak{S}', \mathfrak{g}', v' \rangle$ and a truth value $t \in \mathcal{H}$ ($t \neq \bot$), a non-empty relation $Z \subseteq \mathfrak{S} \times \mathfrak{S}'$ is a t-bisimulation between $\mathfrak{M}$ and $\mathfrak{M}'$ if for any pair $\langle \mathfrak{s}, \mathfrak{s}' \rangle \in Z$

(*base*) $t \wedge v(\mathfrak{s}, P) = t \wedge v'(\mathfrak{s}', P)$ for every $P \in \Phi$

(*forth*) for every $\mathfrak{r} \in \mathfrak{S}$ such that $t \wedge \mathfrak{g}(\mathfrak{s}, \mathfrak{r}) \neq \bot$, there exists an $\mathfrak{r}' \in \mathfrak{S}'$ such that $t \wedge \mathfrak{g}(\mathfrak{s}, \mathfrak{r}) = t \wedge \mathfrak{g}'(\mathfrak{s}', \mathfrak{r}')$ and $\mathfrak{r} Z \mathfrak{r}'$

(*back*) for every $\mathfrak{r}' \in \mathfrak{S}'$ such that $t \wedge \mathfrak{g}'(\mathfrak{s}', \mathfrak{r}') \neq \bot$, there exists an $\mathfrak{r} \in \mathfrak{S}$ such that $t \wedge \mathfrak{g}'(\mathfrak{s}', \mathfrak{r}') = t \wedge \mathfrak{g}(\mathfrak{s}, \mathfrak{r})$ and $\mathfrak{r} Z \mathfrak{r}'$

Two states $\mathfrak{s}$ and $\mathfrak{s}'$ are called *t-bisimilar* (notation $\mathfrak{s} \leftrightarroweq_t \mathfrak{s}'$ or $\mathfrak{M}, \mathfrak{s} \leftrightarroweq_t \mathfrak{M}', \mathfrak{s}'$) if there is a t-bisimulation Z between $\mathfrak{M}$ and $\mathfrak{M}'$ such that $\mathfrak{s} Z \mathfrak{s}'$.

We can now prove the basic theorem which states that t-bisimulation implies t-invariance.

Theorem 1. *If $\mathfrak{M}, \mathfrak{s} \leftrightarroweq_t \mathfrak{M}', \mathfrak{s}'$, then $t \wedge v(\mathfrak{s}, X) = t \wedge v'(\mathfrak{s}', X)$ for every $X \in L^{\mathcal{H}}_{\Box\Diamond}$.*

Proof. The proof runs by induction on the formation of X. If $X \in \Phi$, the result follows by the base condition and if $X \in \mathcal{H}$ is a propositional constant it is trivial. In case X is a conjunction $X_1 \wedge X_2$ the result follows by Def. 2 and the idempotency of the meet operation in lattices.

If $X = X_1 \vee X_2$, then

$$\begin{aligned}
t \wedge v(\mathfrak{s}, X_1 \vee X_2) &= t \wedge \Big(v(\mathfrak{s}, X_1) \vee v(\mathfrak{s}, X_2)\Big) && \text{(Def. 2)} \\
&= \Big(t \wedge v(\mathfrak{s}, X_1)\Big) \vee \Big(t \wedge v(\mathfrak{s}, X_2)\Big) && \text{(Distributivity of } \mathcal{H}\text{)} \\
&= \Big(t \wedge v'(\mathfrak{s}', X_1)\Big) \vee \Big(t \wedge v'(\mathfrak{s}', X_2)\Big) && \text{(Inductive Hypothesis)} \\
&= t \wedge \Big(v'(\mathfrak{s}', X_1) \vee v'(\mathfrak{s}', X_2)\Big) \\
&= t \wedge v'(\mathfrak{s}', X_1 \vee X_2)
\end{aligned}$$

If $X = X_1 \supset X_2$, then

$$\begin{aligned}
t \wedge v(\mathfrak{s}, X_1 \supset X_2) &= t \wedge \Big(v(\mathfrak{s}, X_1) \Rightarrow v(\mathfrak{s}, X_2)\Big) \\
&= t \wedge \Bigg(\Big(t \wedge v(\mathfrak{s}, X_1)\Big) \Rightarrow \Big(t \wedge v(\mathfrak{s}, X_2)\Big)\Bigg) && \text{(Lemma 1(4))} \\
&= t \wedge \Bigg(\Big(t \wedge v'(\mathfrak{s}', X_1)\Big) \Rightarrow \Big(t \wedge v'(\mathfrak{s}', X_2)\Big)\Bigg) && \text{(Inductive Hypothesis)} \\
&= t \wedge \Big(v'(\mathfrak{s}', X_1) \Rightarrow v'(\mathfrak{s}', X_2)\Big) && \text{(Lemma 1(4))} \\
&= t \wedge v'(\mathfrak{s}', X_1 \supset X_2)
\end{aligned}$$

For the case $X = \Box X_1$, the proof is split in two inequalities. Let $\mathfrak{S}_t = \{\mathfrak{r} \in \mathfrak{S} \mid t \wedge \mathfrak{g}(\mathfrak{s}, \mathfrak{r}) \neq \bot\}$ and $\mathfrak{S}'_t$ be the set that contains all the states $\mathfrak{r}'$ specified in the forth condition. The fourth step below, where $\mathfrak{r}$ is restricted to range over $\mathfrak{S}_t$ instead of $\mathfrak{S}$, is justified by Lemma 1(2).

$$
\begin{aligned}
t \wedge v(\mathfrak{s}, \Box X_1) &= t \wedge \bigwedge_{\mathfrak{r} \in \mathfrak{S}} \Big(\mathfrak{g}(\mathfrak{s}, \mathfrak{r}) \Rightarrow v(\mathfrak{r}, X_1)\Big) \\
&= \bigwedge_{\mathfrak{r} \in \mathfrak{S}} \Big(t \wedge \Big(\mathfrak{g}(\mathfrak{s}, \mathfrak{r}) \Rightarrow v(\mathfrak{r}, X_1)\Big)\Big) \\
&= \bigwedge_{\mathfrak{r} \in \mathfrak{S}} t \wedge \Big((t \wedge \mathfrak{g}(\mathfrak{s}, \mathfrak{r})) \Rightarrow (t \wedge v(\mathfrak{r}, X_1))\Big) && \text{(Lemma 1(4))} \\
&= \bigwedge_{\mathfrak{r} \in \mathfrak{S}_t} t \wedge \Big((t \wedge \mathfrak{g}(\mathfrak{s}, \mathfrak{r})) \Rightarrow (t \wedge v(\mathfrak{r}, X_1))\Big) && (\mathfrak{r} \notin \mathfrak{S}_t \text{ does not affect}) \\
&= \bigwedge_{\mathfrak{r}' \in \mathfrak{S}'_t} t \wedge \Big(\underbrace{t \wedge \mathfrak{g}'(\mathfrak{s}', \mathfrak{r}')}_{\text{forth}} \Rightarrow \underbrace{t \wedge v'(\mathfrak{r}', X_1)}_{\text{Induct. Hypothesis}}\Big) && \text{(Lemma 1(3), Def. 4)} \\
&= t \wedge \bigwedge_{\mathfrak{r}' \in \mathfrak{S}'_t} \Big(\mathfrak{g}'(\mathfrak{s}', \mathfrak{r}') \Rightarrow v'(\mathfrak{r}', X_1)\Big) && \text{(Lemma 1(4))} \\
&\geq t \wedge \bigwedge_{\mathfrak{r}' \in \mathfrak{S}'} \Big(\mathfrak{g}'(\mathfrak{s}', \mathfrak{r}') \Rightarrow v'(\mathfrak{r}', X_1)\Big) && (\mathfrak{S}'_t \subseteq \mathfrak{S}') \\
&= t \wedge v'(\mathfrak{s}', \Box X_1)
\end{aligned}
$$

The proof of the other inequality ($\leq$) runs in a completely symmetric way by the back condition. For the case of the other modal operator, we provide below the argument for the first direction ($\leq$):

$$
\begin{aligned}
t \wedge v(\mathfrak{s}, \Diamond X_1) &= t \wedge \bigvee_{\mathfrak{r} \in \mathfrak{S}} \Big(\mathfrak{g}(\mathfrak{s},\mathfrak{r}) \wedge v(\mathfrak{r}, X_1)\Big) \\
&= \bigvee_{\mathfrak{r} \in \mathfrak{S}} \Big(t \wedge \mathfrak{g}(\mathfrak{s},\mathfrak{r}) \wedge v(\mathfrak{r}, X_1)\Big) && \text{(Lemma 1(5))} \\
&= \bigvee_{\mathfrak{r} \in \mathfrak{S}} \Big(\Big(t \wedge \mathfrak{g}(\mathfrak{s},\mathfrak{r})\Big) \wedge \Big(t \wedge v(\mathfrak{r}, X_1)\Big)\Big) \\
&= \bigvee_{\mathfrak{r} \in \mathfrak{S}_t} \Big(\Big(t \wedge \mathfrak{g}(\mathfrak{s},\mathfrak{r})\Big) \wedge \Big(t \wedge v(\mathfrak{r}, X_1)\Big)\Big) \\
&= \bigvee_{\mathfrak{r}' \in \mathfrak{S}'_t} \Big(\Big(\underbrace{t \wedge \mathfrak{g}'(\mathfrak{s}',\mathfrak{r}')}_{\text{forth}}\Big) \wedge \Big(\underbrace{t \wedge v'(\mathfrak{r}', X_1)}_{\text{Induct. Hypothesis}}\Big)\Big) \\
&\leq \bigvee_{\mathfrak{r}' \in \mathfrak{S}'} \Big(\Big(t \wedge \mathfrak{g}'(\mathfrak{s}',\mathfrak{r}'))\Big) \wedge \Big(t \wedge v'(\mathfrak{r}', X_1)\Big)\Big) && (\mathfrak{S}'_t \subseteq \mathfrak{S}') \\
&= t \wedge \bigvee_{\mathfrak{r}' \in \mathfrak{S}'} \Big(\mathfrak{g}'(\mathfrak{s}',\mathfrak{r}')) \wedge v'(\mathfrak{r}', X_1)\Big) && \text{(Lemma 1(5))} \\
&= t \wedge v'\Big(\mathfrak{s}', \Diamond X_1\Big)
\end{aligned}
$$

The other direction is symmetric and the proof of the theorem is complete. ∎

EF-type games for t-bisimulation The t-bisimulation game is a simple variant of the Ehrenfeucht-Fraïssé game played in First-Order Logic. For the purposes of the rest of this section call a state $\mathfrak{r}$ a *t-compatible successor* state of $\mathfrak{s}$ if $t \wedge \mathfrak{g}(\mathfrak{s},\mathfrak{r}) \neq \bot$. Two elements a, a' of $\mathcal{H}$ are called *$t_\wedge$-equivalent* if $t \wedge a = t \wedge a'$. We call *labels* the $\mathcal{H}$-truth values attached to the graph's edges and to the propositional letters of the language in each possible world. The t-bisimulation game is played on two *pointed* $\mathcal{H}$-models (models with a single distinguished state) $\mathfrak{M}, \mathfrak{s}_0$ and $\mathfrak{M}', \mathfrak{s}'_0$. There exists one marked element in each $\mathcal{H}$-model; initially, the marked elements are the distinguished nodes $\mathfrak{s}_0$ and $\mathfrak{s}'_0$. In each *round* of the game

- **Player I** selects one of the $\mathcal{H}$-models, chooses a t-compatible successor of the marked element and moves the marker along the edge (labelled by a_I) to its target
- **Player II** responds with a move of the marker in the other $\mathcal{H}$-model in a corresponding t-compatible transition (labelled by a_{II}) such that a_I and a_{II} are $t_\wedge$-equivalent and the labels of the propositional letters in the marked elements (states) of the models are also $t_\wedge$-equivalent

The *length* of the game is the (finite or infinite) number of rounds and Player II loses the *match* if at a certain round cannot respond with an appropriate move.

It is obvious that Player I is trying to spoil a t-bisimulation while Player II is trying to reveal one. Player II has a *winning strategy* in a game of n rounds if she can win every n-round game played on $\mathfrak{M}, \mathfrak{s}_0$ and $\mathfrak{M}', \mathfrak{s}'_0$. In a classical fashion, we can proceed to a finer analysis of t-bisimulations using the inductively defined notion of the modal depth of a formula (the maximum number of modal operators encountered in a subformula, [BdRV01,MV03]). The notion of a *t-bisimulation bounded by a positive integer n*, or any ordinal actually, can easily be defined (see [Ger99, Chapter 2.1] for the classical case), but we will not give further details here, since the whole construction is identical to the classical one. We only provide the following proposition which generalizes the known classical results from two-valued modal logic:

Proposition 1. *1. Player II has a winning strategy in the n-round game played on $\mathfrak{M}, \mathfrak{s}_0$ and $\mathfrak{M}', \mathfrak{s}'_0$ iff modal truth is t-invariant in $\mathfrak{s}_0$ and $\mathfrak{s}'_0$ for every formula up to modal depth n.*
2. Player II has a winning strategy for the infinite game played on $\mathfrak{M}, \mathfrak{s}_0$ and $\mathfrak{M}', \mathfrak{s}'_0$ iff $\mathfrak{s}_0 \leftrightarrows_t \mathfrak{s}'_0$.

Proof. The proof of the first item runs by induction on n and is actually a restatement of the proof of Theorem 1. The second item follows by the definitions above. ∎

The t-bisimulation games can be formulated in a simple way for the class of languages built on finite linear orders. Assuming further that truth values are colours, linearly ordered, and given that the meet operation is simple in finite chains ($a \wedge b = min(a, b)$) the game can be described in an easy way that provides also an element of fun.

t-unravellings and the tree-model property The idea of unravelling a model into a modally-equivalent tree model is known both from modal logic and the theory of processes. In the latter field, the states of the unravelled model represent *traces* (*histories*) of processes, starting from a state $\mathfrak{s}$. The following definition provides the many-valued analog of this notion.

Definition 5. Given a pointed model $\mathfrak{M} = \langle \mathfrak{S}, \mathfrak{g}, v \rangle, \mathfrak{s}_1$, its *$t$-unravelling* is the model $\mathfrak{M}^u_{\mathfrak{s}_1} = \langle \mathfrak{S}^u_{\mathfrak{s}_1}, \mathfrak{g}^u_{\mathfrak{s}_1}, v^u_{\mathfrak{s}_1} \rangle$, where

1. $\mathfrak{S}^u_{\mathfrak{s}_1}$ consists of all tuples $\langle \mathfrak{s}_1, \cdots, \mathfrak{s}_k \rangle$ where $\mathfrak{s}_{i+1}$ is a t-compatible successor of $\mathfrak{s}_i$
2. $\mathfrak{g}^u_{\mathfrak{s}_1}(\langle \mathfrak{s}_1, \cdots, \mathfrak{s}_k \rangle, \langle \mathfrak{s}_1, \cdots, \mathfrak{s}_{k+1} \rangle) = t \wedge \mathfrak{g}(\mathfrak{s}_k, \mathfrak{s}_{k+1})$, and $\perp$ for any other pair of tuples
3. $v^u_{\mathfrak{s}_1}(\langle \mathfrak{s}_1, \cdots, \mathfrak{s}_k \rangle, P) = t \wedge v(\mathfrak{s}_k, P)$, $(P \in \Phi)$

Obviously $\mathfrak{M}^u_{\mathfrak{s}_1}$ is a tree model and the following proposition can be proved by a careful inspection on the definition of a t-bisimulation.

Proposition 2. The graph of the function from $\mathfrak{S}^u_{\mathfrak{s}_1}$ to $\mathfrak{S}$, which maps every tuple to its last component (and $\langle \mathfrak{s}_1 \rangle$ to $\mathfrak{s}_1$) is a t-bisimulation.

Thus, modal truth is t-invariant for the transition from $\mathfrak{s}_1$ to the root $\langle \mathfrak{s}_1 \rangle$ of the tree and this is a generalized version of the *tree model property* [Var97].

Satisfiability in Many-Valued Modal Languages The general satisfiability problem in this context can be phrased as follows: given $X \in L^{\mathcal{H}}_{\Box\Diamond}$ and $t \in \mathcal{H}$, is there a state $\mathfrak{s}$ of an $\mathcal{H}$-model $\mathfrak{M}$ in which $v(\mathfrak{s}, X) \geq t$? This is equivalent (by Lemma 1(2)) to $t \Rightarrow v(\mathfrak{s}, X) = \top$ which is equivalent (by Def. 2) to $v(\mathfrak{s}, t \supset X) = \top$. Thus, the general satisfiability problem is subsumed by the question of finding a state in which a formula is mapped to the top element of the lattice. By the previous paragraph, if such a state/model exists, then this formula can be also satisfied at the root of a ($\top$-unravelled) tree. Imitating the classical arguments ([MV03]), it is easy to prove that, if $\mathcal{H}$ is finite, every formula can be satisfied in a finite tree whose size is bounded: its depth is bounded by the modal depth of X and its branching degree is bounded by the number of box and diamond subformulas of x. This leads to a simple way of proving the fact that the many-valued analog of the system **K** (which is determined by the class of all $\mathcal{H}$-models [Fit92]) has a decidable general satisfiability problem.

3.2 Weak Bisimulations for Many-Valued Modal Languages

We proceed now to define, a weaker, more fine-grained notion of bisimulation that is directly inspired from (and can be better explained in the context of) the multiple-expert semantics of these languages. We first fix some notation.
Let $I_{\mathcal{H}}$ denote the set of join-irreducible elements of $\mathcal{H}$. For the rest of this section, we fix a complete Heyting algebra $\mathcal{H}$ that possesses the following property:

$$\text{Every } t \in \mathcal{H} - I_{\mathcal{H}} \text{ is equal to the join of a finite number of elements in } I_{\mathcal{H}}. \tag{1}$$

Define the function $D_{\mathcal{H}}$ from $\mathcal{H} - \{\bot\}$ to $2^{I_{\mathcal{H}}}$, such that $D_{\mathcal{H}}(t) = \{c \in I_{\mathcal{H}} \mid c \leq t\}$. Using Property (1), we see that $t = \bigvee_{c \in D_{\mathcal{H}}(t)} c$. Intuitively, $D_{\mathcal{H}}$ provides a decomposition of a value $t \in \mathcal{H}$ into join-irreducible values. In the next definition, a bisimulation relation is defined for every truth value, but in a way that it is "upwards (with respect to the lattice of truth values) consistent".

Definition 6 (Weak bisimulation). Given two $\mathcal{H}$-models $\mathfrak{M} = \langle \mathfrak{S}, \mathfrak{g}, v\rangle$ and $\mathfrak{M}' = \langle \mathfrak{S}', \mathfrak{g}', v'\rangle$, a function Z from $\mathcal{H} - \{\bot\}$ to $2^{\mathfrak{S} \times \mathfrak{S}'}$ is a *weak bisimulation between* $\mathfrak{M}$ *and* $\mathfrak{M}'$ if it satisfies the following properties*:*

- for every $t_1, t_2 \in \mathcal{H}$
  (*consistency*) $Z(t_1 \vee t_2) = Z(t_1) \cap Z(t_2)$
- for every join-irreducible value $t \in I_{\mathcal{H}}$ and any pair $\langle \mathfrak{s}, \mathfrak{s}'\rangle \in Z(t)$
  (*base*) $t \wedge v(\mathfrak{s}, P) = t \wedge v'(\mathfrak{s}', P)$ for every $P \in \Phi$
  (*forth*) for every $\mathfrak{r} \in \mathfrak{S}$ such that $t \wedge \mathfrak{g}(\mathfrak{s}, \mathfrak{r}) \neq \bot$
  and for every $c \in D_{\mathcal{H}}(t \wedge \mathfrak{g}(\mathfrak{s}, \mathfrak{r}))$,
  there exists an $\mathfrak{r}' \in \mathfrak{S}'$ such that $c \leq \mathfrak{g}'(\mathfrak{s}', \mathfrak{r}')$ and $\langle \mathfrak{r}, \mathfrak{r}'\rangle \in Z(c)$
  (*back*) for every $\mathfrak{r}' \in \mathfrak{S}'$ such that $t \wedge \mathfrak{g}'(\mathfrak{s}', \mathfrak{r}') \neq \bot$
  and for every $c \in D_{\mathcal{H}}(t \wedge \mathfrak{g}'(\mathfrak{s}', \mathfrak{r}'))$,
  there exists an $\mathfrak{r} \in \mathfrak{S}$ such that $c \leq \mathfrak{g}(\mathfrak{s}, \mathfrak{r})$ and $\langle \mathfrak{r}, \mathfrak{r}'\rangle \in Z(c)$

Two states $\mathfrak{s}$ and $\mathfrak{s}'$ are called *weakly t-bisimilar* (notation $\mathfrak{s} \leftrightsquigarrow_t \mathfrak{s}'$ or $\mathfrak{M}, \mathfrak{s} \leftrightsquigarrow_t \mathfrak{M}', \mathfrak{s}'$) if there is a weak bisimulation Z between $\mathfrak{M}$ and $\mathfrak{M}'$ such that $\langle \mathfrak{s}, \mathfrak{s}' \rangle$ belongs to $Z(t)$.

The reader can check that we have indeed defined a weaker notion than that of a t-bisimulation:

Remark 1. $\mathfrak{M}, \mathfrak{s} \leftrightarrows_t \mathfrak{M}', \mathfrak{s}'$ implies $\mathfrak{M}, \mathfrak{s} \leftrightsquigarrow_t \mathfrak{M}', \mathfrak{s}'$.

The basic theorem of the previous section is still valid under this new notion, but the proof requires some more elaboration.

Theorem 2. *If* $\mathfrak{M}, \mathfrak{s} \leftrightsquigarrow_t \mathfrak{M}', \mathfrak{s}'$*, then* $t \wedge v(\mathfrak{s}, X) = t \wedge v'(\mathfrak{s}', X)$ *for every* $X \in L^{\mathcal{H}}_{\Box\Diamond}$.

Proof. We first prove that the theorem holds in the case that t is join-irreducible. The proof runs by induction on the formation of X. If $X \in \Phi$, the result follows by the base condition and if $X \in \mathcal{H}$ is a propositional constant it is trivial. For the cases in which X is a formula of the form $X_1 \wedge X_2$, $X_1 \vee X_2$ or $X_1 \supset X_2$, the result can be obtaind as in the proof Theorem 1.

For the case $X = \Box X_1$, the proof is split in two inequalities. Let $\mathfrak{S}_t = \{\mathfrak{r} \in \mathfrak{S} \mid t \wedge \mathfrak{g}(\mathfrak{s}, \mathfrak{r}) \neq \bot\}$. The fourth step below, where $\mathfrak{r}$ is restricted to range over $\mathfrak{S}_t$

instead of $\mathfrak{S}$, is justified by Lemma 1(2). We denote by $\mathfrak{R}'$ the set that contains all the states $\mathfrak{r}'$ specified in the definition of the forth condition.

$$
\begin{aligned}
t \wedge v(\mathfrak{s}, \Box X_1) &= t \wedge \bigwedge_{\mathfrak{r} \in \mathfrak{S}} \Big(\mathfrak{g}(\mathfrak{s},\mathfrak{r}) \Rightarrow v(\mathfrak{r}, X_1) \Big) \\
&= \bigwedge_{\mathfrak{r} \in \mathfrak{S}} \Big(t \wedge \Big(\mathfrak{g}(\mathfrak{s},\mathfrak{r}) \Rightarrow v(\mathfrak{r}, X_1) \Big) \Big) \\
&= \bigwedge_{\mathfrak{r} \in \mathfrak{S}} \Big(t \wedge \Big((t \wedge \mathfrak{g}(\mathfrak{s},\mathfrak{r})) \Rightarrow (t \wedge v(\mathfrak{r}, X_1)) \Big) \Big) && \text{(Lemma 1(4))} \\
&= \bigwedge_{\mathfrak{r} \in \mathfrak{S}_t} \Big(t \wedge \Big((t \wedge \mathfrak{g}(\mathfrak{s},\mathfrak{r})) \Rightarrow (t \wedge v(\mathfrak{r}, X_1)) \Big) \Big) && (\mathfrak{r} \notin \mathfrak{S}_t \text{ does not affect}) \\
&= \bigwedge_{\mathfrak{r} \in \mathfrak{S}_t} \Big(t \wedge \Big((\bigvee_{c \in D_{\mathcal{H}}(t \wedge \mathfrak{g}(\mathfrak{s},\mathfrak{r}))} c) \Rightarrow (t \wedge v(\mathfrak{r}, X_1)) \Big) \Big) && \text{(Def. of } D_{\mathcal{H}}) \\
&= \bigwedge_{\mathfrak{r} \in \mathfrak{S}_t} \Big(t \wedge \bigwedge_{c \in D_{\mathcal{H}}(t \wedge \mathfrak{g}(\mathfrak{s},\mathfrak{r}))} \Big(c \Rightarrow (t \wedge v(\mathfrak{r}, X_1)) \Big) \Big) && \text{(Lemma 1(6))} \\
&= \bigwedge_{\mathfrak{r} \in \mathfrak{S}_t} \; \bigwedge_{c \in D_{\mathcal{H}}(t \wedge \mathfrak{g}(\mathfrak{s},\mathfrak{r}))} \Big(t \wedge \Big((t \wedge c) \Rightarrow (t \wedge v(\mathfrak{r}, X_1)) \Big) \Big) \\
&\geq \bigwedge_{\mathfrak{r}' \in \mathfrak{R}'} \Big(t \wedge \Big((t \wedge \underbrace{\mathfrak{g}'(\mathfrak{s}',\mathfrak{r}'))}_{\text{forth}} \Rightarrow (\underbrace{t \wedge v'(\mathfrak{r}', X_1)}_{\text{Induct. Hypothesis}}) \Big) \Big) && \text{(Lemma 1(3), Def. 4)} \\
&= \bigwedge_{\mathfrak{r}' \in \mathfrak{R}'} \Big(t \wedge \Big(\mathfrak{g}'(\mathfrak{s}',\mathfrak{r}') \Rightarrow v'(\mathfrak{r}', X_1) \Big) \Big) && \text{(Lemma 1(4))} \\
&= t \wedge \bigwedge_{\mathfrak{r}' \in \mathfrak{R}'} \Big(\mathfrak{g}'(\mathfrak{s}',\mathfrak{r}') \Rightarrow v'(\mathfrak{r}', X_1) \Big) \\
&\geq t \wedge \bigwedge_{\mathfrak{r}' \in \mathfrak{S}'} \Big(\mathfrak{g}'(\mathfrak{s}',\mathfrak{r}') \Rightarrow v'(\mathfrak{r}', X_1) \Big) && (\mathfrak{R}' \subseteq \mathfrak{S}') \\
&= t \wedge v'(\mathfrak{s}', \Box X_1)
\end{aligned}
$$

The proof of the other inequality ($\leq$) runs in a completely symmetric way by the back condition. For the case of the other modal operator, we provide below the argument for the first direction ($\leq$):

$$
\begin{aligned}
t \wedge v(\mathfrak{s}, \Diamond X_1) &= t \wedge \bigvee_{\mathfrak{r} \in \mathfrak{S}} \Big(\mathfrak{g}(\mathfrak{s},\mathfrak{r}) \wedge v(\mathfrak{r}, X_1)\Big) \\
&= \bigvee_{\mathfrak{r} \in \mathfrak{S}} \Big(t \wedge \Big(\mathfrak{g}(\mathfrak{s},\mathfrak{r}) \wedge v(\mathfrak{r}, X_1)\Big)\Big) && \text{(Distributivity)} \\
&= \bigvee_{\mathfrak{r} \in \mathfrak{S}} \Big((t \wedge \mathfrak{g}(\mathfrak{s},\mathfrak{r})) \wedge (t \wedge v(\mathfrak{r}, X_1))\Big) \\
&= \bigvee_{\mathfrak{r} \in \mathfrak{S}_t} \Big((t \wedge \mathfrak{g}(\mathfrak{s},\mathfrak{r})) \wedge (t \wedge v(\mathfrak{r}, X_1))\Big) && (\mathfrak{r} \notin \mathfrak{S}_t \text{ does not affect}) \\
&= \bigvee_{\mathfrak{r} \in \mathfrak{S}_t} \Big(\big(\bigvee_{c \in D_{\mathcal{H}}(t \wedge \mathfrak{g}(\mathfrak{s},\mathfrak{r}))} c\big) \wedge (t \wedge v(\mathfrak{r}, X_1))\Big) && \text{(Def. of } D_{\mathcal{H}}) \\
&= \bigvee_{\mathfrak{r} \in \mathfrak{S}_t} \bigvee_{c \in D_{\mathcal{H}}(t \wedge \mathfrak{g}(\mathfrak{s},\mathfrak{r}))} \Big(c \wedge (t \wedge v(\mathfrak{r}, X_1))\Big) && \text{(Distributivity)} \\
&= \bigvee_{\mathfrak{r} \in \mathfrak{S}_t} \bigvee_{c \in D_{\mathcal{H}}(t \wedge \mathfrak{g}(\mathfrak{s},\mathfrak{r}))} \Big((t \wedge c) \wedge (t \wedge v(\mathfrak{r}, X_1))\Big) \\
&\leq \bigvee_{\mathfrak{r}' \in \mathfrak{R}'} \Big((t \wedge \underbrace{\mathfrak{g}'(\mathfrak{s}',\mathfrak{r}')}_{\text{forth}}) \wedge (\underbrace{t \wedge v'(\mathfrak{r}', X_1)}_{\text{Induct. Hypothesis}})\Big) && \text{(Monotonicity of } \wedge\text{, Def. 4)} \\
&= \bigvee_{\mathfrak{r}' \in \mathfrak{R}'} \Big(t \wedge \Big(\mathfrak{g}'(\mathfrak{s}',\mathfrak{r}') \wedge v'(\mathfrak{r}', X_1)\Big)\Big) \\
&= t \wedge \bigvee_{\mathfrak{r}' \in \mathfrak{R}'} \Big(\mathfrak{g}'(\mathfrak{s}',\mathfrak{r}') \wedge v'(\mathfrak{r}', X_1)\Big) && \text{(Distributivity)} \\
&\leq t \wedge \bigvee_{\mathfrak{r}' \in \mathfrak{S}'} \Big(\mathfrak{g}'(\mathfrak{s}',\mathfrak{r}') \wedge v'(\mathfrak{r}', X_1)\Big) && (\mathfrak{R}' \subseteq \mathfrak{S}') \\
&= t \wedge v'(\mathfrak{s}', \Diamond X_1)
\end{aligned}
$$

The other direction is symmetric and the proof for that case in which t is join-irreducible is complete.

Suppose now that t is not join-irreducible. Then:

$$
\begin{aligned}
t \wedge v(\mathfrak{s}, X) &= \Big(\bigvee_{c \in D_{\mathcal{H}}(t)} c\Big) \wedge v(\mathfrak{s}, X) && \text{(Def. of } D_{\mathcal{H}}\text{)}\\
&= \bigvee_{c \in D_{\mathcal{H}}(t)} \Big(c \wedge v(\mathfrak{s}, X)\Big) && \text{(Distributivity)}\\
&= \bigvee_{c \in D_{\mathcal{H}}(t)} \Big(c \wedge v'(\mathfrak{s}', X)\Big) && (c \text{ is join-irreducible})\\
&= \Big(\bigvee_{c \in D_{\mathcal{H}}(t)} c\Big) \wedge v'(\mathfrak{s}', X) && \text{(Distributivity)}\\
&= t \wedge v'(\mathfrak{s}', X) && \text{(Def. of } D_{\mathcal{H}}\text{)}
\end{aligned}
$$

∎

Image-finite $\mathcal{H}$-models and weak t-bisimulations One of the fundamental questions in the bisimulation-based analysis of modal languages, concerns the identification of cases in which the converse of Theorem 2 is true. Much obviously, it is not always true: the classical counterexample of two tree models, both with a finite branch for each natural number, one of which possesses an infinite branch, suffices (cf. [BdRV01, Chapter 2.2]). The simplest example of Hennessy-Milner classes of modal models (classes in which modal equivalence is itself a bisimulation relation) is the class of *image-finite* models, in which each state has only a finite number of successors. It is natural to consider a straightforward many-valued analog of this notion by considering $\mathcal{H}$-models in which for each state $\mathfrak{s}$, the set of successors of $\mathfrak{s}$ is always finite and check whether in this case t-invariance implies t-bisimilarity. Formally, the notion of image-finite $\mathcal{H}$-models is defined as follows.

Definition 7 (Image-finite $\mathcal{H}$-models). An $\mathcal{H}$-model $\mathfrak{M} = \langle \mathfrak{S}, \mathfrak{g}, v\rangle$ is called *image-finite* if for every $\mathfrak{s} \in \mathfrak{S}$, the set $S_{\mathfrak{s}} = \{\mathfrak{r} \in \mathfrak{S} \mid \mathfrak{g}(\mathfrak{s}, \mathfrak{r}) \neq \bot\}$ is finite.

The following theorem states that for image-finite $\mathcal{H}$-models, t-invariance implies t-bisimilarity.

Theorem 3. *Let $\mathfrak{M} = \langle \mathfrak{S}, \mathfrak{g}, v\rangle$ and $\mathfrak{M}' = \langle \mathfrak{S}', \mathfrak{g}', v'\rangle$ be image-finite $\mathcal{H}$-models for $L^{\mathcal{H}}_{\Box\Diamond}(\Phi)$. Define the function Z from $\mathcal{H} - \{\bot\}$ to $2^{\mathfrak{S}\times\mathfrak{S}'}$ so that for every $\mathfrak{s} \in \mathfrak{S}$, every $\mathfrak{s}' \in \mathfrak{S}'$ and every $t \in \mathcal{H} - \{\bot\}$, $\langle \mathfrak{s}, \mathfrak{s}'\rangle \in Z(t)$ iff modal truth is t-invariant for the transition between $\mathfrak{s}$ and $\mathfrak{s}'$. Then Z is a weak bisimulation between $\mathfrak{M}$ and $\mathfrak{M}'$.*

Proof. The consistency condition is easy to verify, and the base condition staightforward. We next show that the Z satisfies the forth condition; the proof for the back condition is completely symmetric.

Suppose for the sake of contradiction that Z does not satisfy the forth condition. This means that there exist a truth value $t \in I_{\mathcal{H}}$, a pair $\langle \mathfrak{s}, \mathfrak{s}'\rangle \in Z(t)$, a state $\mathfrak{r} \in \mathfrak{S}$ such that $t \wedge \mathfrak{g}(\mathfrak{s}, \mathfrak{r}) \neq \bot$ and a truth value $c \in D_{\mathcal{H}}(t \wedge \mathfrak{g}(\mathfrak{s}, \mathfrak{r}))$, such that for every $\mathfrak{r}' \in \mathfrak{S}'$, if $c \leq \mathfrak{g}'(\mathfrak{s}', \mathfrak{r}')$ then $\langle \mathfrak{r}, \mathfrak{r}' \rangle \notin Z(c)$.

Since $\mathfrak{M}'$ is image finite, the set $R' = \{\mathfrak{r}' \in \mathfrak{S}' \mid c \leq \mathfrak{g}'(\mathfrak{s}', \mathfrak{r}')\}$ is finite. We first show that R' is non-empty. We have:

$$\begin{aligned} c \leq t \wedge \mathfrak{g}(\mathfrak{s}, \mathfrak{r}) \leq t \wedge \bigvee_{\mathfrak{r} \in \mathfrak{S}} \big(\mathfrak{g}(\mathfrak{s}, \mathfrak{r}) \wedge \top\big) &= t \wedge v(\mathfrak{s}, \top) = t \wedge v'(\mathfrak{s}', \top) \\ &= t \bigvee_{\mathfrak{r}' \in \mathfrak{S}'} \big(\mathfrak{g}'(\mathfrak{s}', \mathfrak{r}') \wedge \top\big) = \bigvee_{\mathfrak{r}' \in \mathfrak{S}'} \big(t \wedge \mathfrak{g}'(\mathfrak{s}', \mathfrak{r}')\big) \end{aligned}$$

Since $c \in I_{\mathcal{H}}$ and $\mathcal{H}$ is distributive, one can show that for some $\mathfrak{r}' \in \mathfrak{S}'$, $c \leq t \wedge \mathfrak{g}'(\mathfrak{s}', \mathfrak{r}')$, and thus $c \leq \mathfrak{g}'(\mathfrak{s}', \mathfrak{r}')$.

Suppose that $R' = \{\mathfrak{r}'_1, \mathfrak{r}'_2, \ldots, \mathfrak{r}'_k\}$. Then for every i, $1 \leq i \leq k$, $\langle \mathfrak{r}, \mathfrak{r_i}' \rangle \notin Z(c)$, which implies that there exists a formula X_i such that $c \wedge v(\mathfrak{r}, X_i) \neq c \wedge v'(\mathfrak{r_i}', X_i)$. We will define a new formula Y_i such that $c \wedge v(\mathfrak{r}, Y_i) = c$ and $c \wedge v'(\mathfrak{r_i}', Y_i) < c$. Let $a_i = c \wedge v(\mathfrak{r}, X_i)$ and $b_i = c \wedge v'(\mathfrak{r_i}', X_i)$. We consider two cases:

Case 1: $a_i \leq b_i$. Since $a_i \neq b_i$ it is $b_i \not\leq a_i$. Define $Y_i = X_i \supset a_i$. Then, $c \wedge v(\mathfrak{r}, Y_i) = c \wedge v(\mathfrak{r}, X_i \supset a_i) = c \wedge (v(\mathfrak{r}, X_i) \Rightarrow a_i) = c \wedge (a_i \Rightarrow a_i) = c$. Moreover, $c \wedge v'(\mathfrak{r_i}', Y_i) = c \wedge v'(\mathfrak{r_i}', X_i \supset a_i) = c \wedge (v'(\mathfrak{r_i}', X_i) \Rightarrow a_i) = c \wedge (b_i \Rightarrow a_i)$. If $c \wedge (b_i \Rightarrow a_i) = c$, then $c \leq (b_i \Rightarrow a_i)$, which implies $b_i = c \wedge b_i \leq a_i$ (contradiction). Therefore, $c \wedge v'(\mathfrak{r_i}', Y_i) \neq c$.

Case 2: $a_i \not\leq b_i$. Define $Y_i = a_i \supset X_i$. Then, $c \wedge v(\mathfrak{r}, Y_i) = c \wedge v(\mathfrak{r}, a_i \supset X_i) = c \wedge (a_i \Rightarrow v(\mathfrak{r}, X_i)) = c \wedge (a_i \Rightarrow a_i) = c$. Moreover, $c \wedge v'(\mathfrak{r_i}', Y_i) = c \wedge v'(\mathfrak{r_i}', a_i \supset X_i) = c \wedge (a_i \Rightarrow v'(\mathfrak{r_i}', X_i)) = c \wedge (a_i \Rightarrow b_i)$. If $c \wedge (a_i \Rightarrow b_i) = c$, then $c \leq (a_i \Rightarrow b_i)$, which implies $a_i = c \wedge a_i \leq b_i$ (contradiction). Therefore, $c \wedge v'(\mathfrak{r_i}', Y_i) \neq c$.

Let $Y = Y_1 \wedge Y_2 \wedge \cdots \wedge Y_k$. It is easy to see that we again have $c \wedge v(\mathfrak{r}, Y) = c$ and $c \wedge v'(\mathfrak{r_i}', Y) < c$.

Let $\phi = \Diamond Y$. Since $\langle \mathfrak{s}, \mathfrak{s}' \rangle \in Z(t)$, it must be $t \wedge v(\mathfrak{s}, \phi) = t \wedge v'(\mathfrak{s}', \phi)$, which implies $c \wedge t \wedge v(\mathfrak{s}, \phi) = c \wedge t \wedge v'(\mathfrak{s}', \phi)$. We compute the two values in the last inequality. For $c \wedge t \wedge v(\mathfrak{s}, \phi)$, we have:

$$\begin{aligned} c \wedge t \wedge v(\mathfrak{s}, \phi) \quad &= c \wedge v(\mathfrak{s}, \phi) \\ &= c \wedge \bigvee_{\mathfrak{u} \in \mathfrak{S}} \Big(\mathfrak{g}(\mathfrak{s}, \mathfrak{u}) \wedge v(\mathfrak{u}, Y\Big) \\ &\geq c \wedge (\mathfrak{g}(\mathfrak{s}, \mathfrak{r}) \wedge v(\mathfrak{r}, Y)) \\ &= (c \wedge \mathfrak{g}(\mathfrak{s}, \mathfrak{r})) \wedge (c \wedge v(\mathfrak{r}, Y)) \\ &= c \wedge c \\ &= c \end{aligned}$$

This implies (since it is also $c \wedge t \wedge v(\mathfrak{s}, \phi) \leq c$) that $c \wedge t \wedge v(\mathfrak{s}, \phi) = c$. For $c \wedge t \wedge v'(\mathfrak{s}', \phi)$, we have:

$$\begin{aligned}
c \wedge t \wedge v'(\mathfrak{s}', \phi) &= c \wedge v'(\mathfrak{s}', \phi) \\
&= c \wedge \bigvee_{\mathfrak{u}' \in \mathfrak{S}'} \Big(\mathfrak{g}'(\mathfrak{s}', \mathfrak{u}') \wedge v'(\mathfrak{u}', Y)\Big) \\
&= \bigvee_{\mathfrak{u}' \in \mathfrak{S}'} \Big(c \wedge \mathfrak{g}'(\mathfrak{s}', \mathfrak{u}') \wedge v'(\mathfrak{u}', Y)\Big) \\
&= \bigvee_{\mathfrak{r}' \in R'} \Big(c \wedge \mathfrak{g}'(\mathfrak{s}', \mathfrak{r}') \wedge v'(\mathfrak{r}', Y)\Big) \vee \bigvee_{\mathfrak{u}' \in \mathfrak{S}' - R'} \Big(c \wedge \mathfrak{g}'(\mathfrak{s}', \mathfrak{u}') \wedge v'(\mathfrak{u}', Y)\Big) \\
&\leq \bigvee_{\mathfrak{r}' \in R'} \Big(c \wedge v'(\mathfrak{r}', Y)\Big) \vee \bigvee_{\mathfrak{u}' \in \mathfrak{S}' - R'} \Big(c \wedge \mathfrak{g}'(\mathfrak{s}', \mathfrak{u}')\Big) \\
&= \bigvee_{\mathfrak{r}' \in R'} \Big(c \wedge v'(\mathfrak{r}', Y)\Big) \vee \bigvee_{\mathfrak{u}' \in \mathfrak{S}' - R'} \Big(c \wedge \mathfrak{g}'(\mathfrak{s}', \mathfrak{u}')\Big) \\
&\leq c
\end{aligned}$$

From $t \wedge c \wedge v(\mathfrak{s}, \phi) = t \wedge c \wedge v'(\mathfrak{s}'\phi)$, we obtain that

$$c = \bigvee_{\mathfrak{r}' \in R'} \Big(c \wedge v'(\mathfrak{r}', Y)\Big) \vee \bigvee_{\mathfrak{u}' \in \mathfrak{S}' - R'} \Big(c \wedge \mathfrak{g}'(\mathfrak{s}', \mathfrak{u}')\Big)$$

However, if $\mathfrak{r}' \in R'$, then $c \wedge v'(\mathfrak{r}', Y) < c$. Furthermore, if $\mathfrak{u}' \in \mathfrak{S}' - R'$, then $c \not\leq \mathfrak{g}'(\mathfrak{s}', \mathfrak{u}')$, which implies that $c \wedge \mathfrak{g}'(\mathfrak{s}', \mathfrak{u}') < c$. By the above equality, c is not join-irreducible, which is a contradiction.
Thus, Z satisfies the forth condition, and the proof is completed. ∎

4 Multiple-Expert Semantics

In this section, we confine ourselves in the class of many-valued modal languages built on *finite* HAs. Each language of this class can be reformulated in a way that is of interest to KR situations involving many interrelated experts. We briefly review below this facet of Fitting's many-valued modal languages. The interested reader should refer to [Fit92, Sect. 1, 3 & 5] for technical details. We then interpret our definitions and results in this alternative context. We note, however, that our results carried out in the previous sections hold in more general Heyting algebras.

A *multiple-expert modal model* is a structure $\langle \mathcal{E}, \mathfrak{S}, \{R_e\}_{e \in \mathcal{E}}, \{v_e\}_{e \in \mathcal{E}}, D \rangle$, such that:

- $\mathcal{E}$ is a finite set of experts.
- D is a partial-order dominance relation on $\mathcal{E}$.
- $\mathfrak{S}$ is a common set of worlds.
- For each $e \in \mathcal{E}$, $\langle \mathfrak{S}, R_e, v_e \rangle$ is a (two-valued) modal model, such that
  (D_1) if $R_e(\mathfrak{s}_1, \mathfrak{s}_2)$ and $D(e, f)$, then $R_f(\mathfrak{s}_1, \mathfrak{s}_2)$, and
  (D_2) for any propositional variable X, if $v_e(\mathfrak{s}, X)$ and $D(e, f)$, then $v_f(\mathfrak{s}, X)$.

The valuations v_e are then properly extended to all modal formulae, so that (D_2) above is preserved.

We are interested now in finding the experts' 'consensus', that is, in elegantly calculating the modal formulae on which our experts agree. This problem can be reformulated as one involving a many-valued language, where sets of experts who agree on the truth of an epistemic statement can be seen as generalized truth values. Note however an important point: by (D_1) and (D_2), **not every set of experts is an 'admissible' generalized truth value**. The 'admissible' sets of experts are those which are dominance-closed, that is, upwards-closed in the order D. The set of all admissible sets of experts form a finite Heyting algebra $\mathcal{H}$ when ordered under set inclusion. We can thus produce an $\mathcal{H}$-model $\langle \mathfrak{S}, \mathfrak{g}, v \rangle$ as follows:

- For $\mathfrak{s}, \mathfrak{r} \in \mathfrak{S}$, $\mathfrak{g}(\mathfrak{s}, \mathfrak{r}) := \{e \in \mathcal{E} : \mathfrak{s} R_e \mathfrak{r}\}$.
- For a propositional variable X, $v(\mathfrak{s}, X) := \{e \in \mathcal{E} : v_e(\mathfrak{s}, X) = 1\}$.

It can then be proved that for any modal formulae ϕ,

$$v(\mathfrak{s}, \phi) = \{e \in \mathcal{E} : v_e(\mathfrak{s}, \phi) = 1\}.$$

We have thus provided a translation of the multiple-expert situation into a many-valued modal model of the language $L^{\mathcal{H}}_{\Box\Diamond}$. The other translation is also feasible. Both translations are presented in [Fit92, Sect. 5].

We are now in the position to express the meaning of our results in this alternative setting. It suffices to observe that in the finite Heyting algebra $\mathcal{H}$ of the 'admissible' subsets of experts, meet is set intersection and join is set union. Assume $\langle \mathfrak{S}, \mathfrak{g}, v \rangle$ and $\langle \mathfrak{S}', \mathfrak{g}', v' \rangle$ are two $\mathcal{H}$-models, and $\in \mathfrak{S}$, $\mathfrak{s}' \in \mathfrak{S}'$. In the case of weak bisimulation we have a sharp description:

- The states $\mathfrak{s}$ and $\mathfrak{s}'$ are weakly t-bisimilar iff for every expert e in t the states $\mathfrak{s}$ and $\mathfrak{s}'$ in the corresponding models $\langle \mathfrak{S}, R_e, v_e \rangle$ and $\langle \mathfrak{S}', R_e, v_e \rangle$ are bisimilar.

An analogous statement *does not* hold for the notion of t-bisimulation, as the example below shows. Note that, together with Remark 1, this implies that the notion of a weak bisimulation is strictly weaker than that of a t-bisimulation.

Example 1. Let $\mathcal{E} = \{e, f\}$ be the set of experts, D empty, $\mathfrak{S} = \{\mathfrak{s}_1, \mathfrak{s}_2, \mathfrak{s}_3\}$ and $\mathfrak{S}' = \{\mathfrak{s}_4, \mathfrak{s}_5\}$ two sets of states, X the unique propositional variable, and:
$R_e(\mathfrak{s}_1, \mathfrak{s}_2)$ and $R_e(\mathfrak{s}_4, \mathfrak{s}_5)$ hold, and R_e fails for any other pair of states,
$R_f(\mathfrak{s}_1, \mathfrak{s}_3)$ and $R_e(\mathfrak{s}_4, \mathfrak{s}_5)$ hold, and R_f fails for any other pair of states,
$v_e(\mathfrak{s}_2, X) = v_e(\mathfrak{s}_5, X) = 1$, and 0 for any other entries,
$v_f(\mathfrak{s}_3, X) = v_f(\mathfrak{s}_5, X) = 1$, and 0 for any other entries.

Then, for the expert e the modal models $\langle \mathfrak{S}, R_e, v_e \rangle$ and $\langle \mathfrak{S}', R_e, v_e \rangle$ are bisimilar, for the expert f the modal models $\langle \mathfrak{S}, R_f, v_f \rangle$ and $\langle \mathfrak{S}', R_f, v_f \rangle$ are bisimilar, but the corresponding $\mathcal{H}$-models $\langle \mathfrak{S}, \mathfrak{g}, v \rangle$ and $\langle \mathfrak{S}', \mathfrak{g}', v' \rangle$ are not $\{e, f\}$-bisimilar.

We can, however, interpret the conditions of the weak t-bisimulation in the multiple-expert scenario as follows:

- The *base* condition of Def. 4 says that moving back and forth between t-bisimilar states does not affect the belief of any expert from the set t, for any propositional letter P.
- The *forth* condition of Def. 4 says that any transition in the second model that involves experts from the fixed set t can be matched with a transition in the first model where all the relevant experts from t are also involved; more experts can also be involved; what we require concerns only those in t.
- Similarly for the *back* condition.

Finally, the meaning of Theorems 1 and 2 is that the bisimulation relation between states of models guarantees the invariance of the epistemic consensus of some experts from a predefined fixed set t. It is also easy to give an equivalent definition of the EF-type games of bisimulation, in terms of the epistemic agreement of the experts.

5 Conclusions - Related Work

In this paper, we have contributed to the extensive literature on the importance and the fundamental nature of bisimulation. Our main aim has been to define a fine-grained notion of bisimulation for Heyting-valued modal languages and establish its basic facts. Our results have an interesting meaning for Knowledge Representation situations, when interpreted in the multiple-expert context. It remains to investigate appropriate extension of smallest and largest bisimulations in this context and address possible applications for Knowledge Engineering in complex epistemic situations.

References

[Acz88] P. Aczel. *Non-Well-Founded Sets.* CSLI Publications, 1988.

[BD74] R. Balbes and Ph. Dwinger. *Distributive Lattices.* University of Missouri Press, 1974.

[BdRV01] P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic.* Number 53 in Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, 2001.

[DP90] B. A. Davey and H. A. Priestley. *Introduction to Lattices and Order.* Cambridge University Press, 1990.

[dR93] M. de Rijke. *Extending Modal Logic.* PhD thesis, Institute for Logic, Language and Computation, University of Amsterdam, 1993.

[EK05] P. Eleftheriou and C. D. Koutras. Frame constructions, truth invariance and validity preservation in many-valued modal logic. *Journal of Applied Non-Classical Logics*, 15(4):367–388, 2005.

[Fit91] M. C. Fitting. Many-valued modal logics. *Fundamenta Informaticae*, 15:235–254, 1991.

[Fit92] M. C. Fitting. Many-valued modal logics II. *Fundamenta Informaticae*, 17:55–73, 1992.

[Ger99] J. D. Gerbrandy. *Bisimulations on Planet Kripke*. PhD thesis, Institute for Logic, Language and Computation, University of Amsterdam, 1999.

[GG84] D. M. Gabbay and F. Guenthner, editors. *Handbook of Philosophical Logic*. D. Reidel, Dordrecht, 1984.

[HM85] M. Hennessy and R. Milner. Algebraic laws for nondeterminism and concurrency. *Journal of the ACM*, 32:137–162, 1985.

[KdR97] N. Kurtonina and M. de Rijke. Bisimulations for temporal logic. *Journal of Logic, Language and Information*, 6:403–425, 1997.

[KNP02] C. D. Koutras, Ch. Nomikos, and P. Peppas. Canonicity and completeness results for many-valued modal logics. *Journal of Applied Non-Classical Logics*, 12(1):7–41, 2002.

[Kou03] C. D. Koutras. A catalog of weak many-valued modal axioms and their corresponding frame classes. *Journal of Applied Non-Classical Logics*, 13(1):47–72, 2003.

[KP02] C. D. Koutras and P. Peppas. Weaker axioms, more ranges. *Fundamenta Informaticae*, 51(3):297–310, 2002.

[MV03] M. Marx and Y. Venema. Local variations on a loose theme: Modal logic and decidability. In M. Y. Vardi and Sc. Weinstein, editors, *Finite Model Theory and its Applications*. Springer Verlag, 2003. To appear.

[Ott99] M. Otto. Bisimulation-invariant Ptime and higher-dimensional mu-calculus. *Theoretical Computer Science*, 224:237–265, 1999.

[Par91] D. Park. Concurrency and automata on infinite sequences. In *Proceedings of the 5th GI Conference*, volume 154 of *Lecture Notes in Computer Science*, pages 561–572. Springer-Verlag, 1991.

[RS70] H. Rasiowa and R. Sikorski. *The Mathematics of Metamathematics*. PWN - Polish Scientific Publishers, Warsaw, third edition, 1970.

[San07] D. Sangiorgi. *On the origins of Bisimulation, Coinduction, and Fixed Points*. Technical Report UBLCS-2007-24, Department of Computer Science, University of Bologna, 2007 (available at: `http://www.cs.unibo.it/~sangio/DOC_public/history_bis_coind.pdf`).

[Var97] M. Y. Vardi. Why is modal logic so robustly decidable? volume 31 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 561–572. AMS, 1997.

[vB83] J. van Benthem. *Modal Logic and Classical Logic*. Bibliopolis, Naples, 1983.

[vB84] J. van Benthem. *Correspondence Theory*, pages 167–247. Volume 2 of Gabbay and Guenthner [GG84], 1984.