

Future Event Logic – Axioms and Complexity

Hans van Ditmarsch

*Logic, University of Sevilla, Spain
hvd@us.es*

Tim French

*Computer Science and Software Engineering, The University of Western Australia
tim@csse.uwa.edu.au*

Sophie Pinchinat

*IRISA, University of Rennes
Sophie.Pinchinat@irisa.fr*

Abstract

In this paper we present a sound and complete axiomatization of future event logic. Future event logic is a logic that generalizes a number of dynamic epistemic logics, by using a new operator \triangleright that acts as a quantifier over the set of all refinements of a given model. (A refinement is like a bisimulation except that from the three relational requirements only ‘atoms’ and ‘back’ need to be satisfied.) Thus the logic combines the simplicity of modal logic with some powers of monadic second order quantification. We prove the axiomatization is sound and complete and discuss some extensions to the result.

Keywords: Bisimulation Quantifier, Modal Logic, Temporal Epistemic Logic, Multi-Agent System

1 Introduction

Modal logic is frequently used for modelling knowledge in multi-agent systems. The semantics of modal logic uses the notion of “possible worlds”, between which an agent is unable to distinguish. In dynamic systems agents acquire new knowledge (say by an announcement, or the execution of some action) that allows to distinguish between worlds they previously could not separate. From the agents point of view, what were “possible worlds” become inconceivable. Thus, a future informative event may be modelled by a

reduction in the agent’s accessibility relation. In [21] the *future event logic* is introduced. It augments the multi-agent logic of knowledge with (only) an operation $\blacktriangleright\phi$ that stands for “ ϕ holds after all informative events” — the diamond version $\blacktriangleright\phi$ stands for “there is an informative event after which ϕ .” The semantics of informative events axiomatized in this paper was presented in [21]; it encompasses action model execution à la Baltag et al [4]: on finite models, it can be easily shown that a model resulting from action model execution is a refinement of the initial model, and for a given refinement of a model we can construct an action model such that its execution is bisimilar to that refinement. Here we examine the important questions that arise for a new logic: expressivity; axiomatizations; and complexity. We visit these questions in both the context of modal logic, and the modal μ -calculus.

Previous works [10,15] have modelled informative events using a notion of model refinement. In [15] it was shown that model restrictions were not sufficient to simulate informative events, and they introduced *refinement trees* for this purpose—a precursor of the semantics of dynamic epistemic logics developed later [22]. We incorporate implicit quantification over informative events directly into the language using a similar notion of *refinement*; in our case a refinement is the inverse of simulation [1]. This work is also closely related to some recent work on bisimulation quantified modal logics [9,11]. The future event operators are weaker operators than bisimulation quantifiers [21], as they are only based on simulations rather than bisimulations, and do not allow us to vary the interpretation of propositional atoms. Bisimulation quantified modal logic has previously been axiomatized by providing a provably correct translation to the modal μ -calculus [8] (albeit a very complicated one).

Thus we may consider *refinement quantification* to be a generalization of future event operators [21] to other modal logics. This is significant in that it motivates the application of the new operator in many different settings: In logics for games [17,2] or in control theory [18,20], it may correspond to a player discarding some moves; for program logics [12] it may correspond to operational refinement [16]; and for topologies it may correspond to sub-space projections.

This paper will present the definitions for refinement quantification in the general settings of modal logic and the modal μ -calculus, and seek to motivate their use in a range of applied logics. We will then address the questions of expressivity, complexity and axiomatization. Specifically: sound and complete axiomatizations will be provided for both modal logic and the modal μ -calculus augmented with refinement quantification; we provide a double exponential upper-bound for each logic; and we show the use of refinement quantification does not change the expressive power of the logics, although they do make each logic exponentially more succinct.

2 Technical preliminaries

Structural notions

Assume a finite set of agents A and a countably infinite set of atoms P .

Definition 2.1 [Structures] A *model* $M = (S, R, V)$ consists of a *domain* S of (factual)

states (or worlds), accessibility $R : A \rightarrow \mathcal{P}(S \times S)$, and a valuation $V : P \rightarrow \mathcal{P}(S)$. For $s \in S$, (M, s) is a state (also known as a pointed Kripke model).

For $R(a)$ we write R_a ; accessibility R can be seen as a set of relations R_a , and V as a set of valuations $V(p)$. Given two states s, s' in the domain, $R_a(s, s')$ means that in state s agent a considers s' a possibility. As we will be often required to discuss several models at once, we will use the convention that $M = (S^M, R^M, V^M)$, $N = (S^N, R^N, V^N)$ etc. Also, given $s \in S^M$, we let M_s refer to the pair (M, s) or the *pointed model*.

In the first instance we will assume that there are no further restrictions on the models. That is, the underlying modal logic is \mathcal{L} whose system of axioms is \mathbf{K} . In future work we will consider how our results may be extended to epistemic logics, such as $\mathbf{S5}$ and $\mathbf{KD45}$.

Definition 2.2 [Bisimulation, simulation, refinement] Let two models $M = (S, R, V)$ and $M' = (S', R', V')$ be given. A non-empty relation $\mathfrak{R} \subseteq S \times S'$ is a bisimulation, iff for all $s \in S$ and $s' \in S'$ with $(s, s') \in \mathfrak{R}$, for all $a \in A$:

atoms $s \in V(p)$ iff $s' \in V'(p)$ for all $p \in P$

forth- a for all $t \in S$, if $R_a(s, t)$, then there is a $t' \in S'$ such that $R'_a(s', t')$ and $(t, t') \in \mathfrak{R}$

back- a for all $t' \in S'$, if $R'_a(s', t')$, then there is a $t \in S$ such that $R_a(s, t)$ and $(t, t') \in \mathfrak{R}$

We write $M_s \Leftrightarrow M'_{s'}$, iff there is a bisimulation between M and M' linking s and s' . Then we call M_s and $M'_{s'}$ bisimilar.

A relation that satisfies **atoms** and **forth- a** for every $a \in A$ is a *simulation*, and in that case $M'_{s'}$ is a *simulation* of M_s , and M_s is a *refinement* of $M'_{s'}$, and we write $M_s \preceq M'_{s'}$ (or $M'_{s'} \succeq M_s$).

A relation that satisfies **atoms** and **forth- b** for every $b \in A$, as well as **back- b** for every $b \in A - \{a\}$ is an *a -simulation*, and in that case $M'_{s'}$ is an *a -simulation* of M_s , and M_s is an *a -refinement* of $M'_{s'}$, and we write $M_s \preceq_a M'_{s'}$ (or $M'_{s'} \succeq_a M_s$).

We note that the definition of simulation and refinement above varies slightly to the one given by Blackburn et al [6]. Here we ensure that simulations and refinements preserve the interpretations of atoms, whereas [6], has them only preserving the truth of atoms. We take this approach as we feel it suits the epistemic domain we aspire to. It is also important to note that in an epistemic setting a refinement corresponds to the *diminishing uncertainty* of agents¹. This means that there is a potential *decrease* in the number of states and transitions in a model. This is perhaps contrary to the concept of program refinement [16] where detail is added to a specification. However, in program refinement the added detail requires a more detailed state space (i.e. extra atoms) and as such is more the domain of bisimulation quantifiers, rather than refinement quantification. It is interesting to note the consequence of program refinement is a more deterministic system which agrees with the notion of diminishing uncertainty.

We give the following lemma for the properties of the relation \succeq_a .

¹ At least, with respect to formulas in which knowledge operators appear within the scope of an even number of negations. It is possible that in a refinement one agent may be less certain about what another agent does not know.

Lemma 2.3 *The relation \succeq_a is reflexive and transitive (a pre-order), and satisfies the Church-Rosser property.*

Proof. Reflexivity follows from the observation that the identity relation satisfies **atoms**, and **back- a** and **forth- a** for all agents a , and therefore also the weaker requirement for refinement. Similarly, given two a -simulations B_1 , and B_2 , we can see that their composition, $\{(x, z) \mid \exists y, (x, y) \in B_1, (y, z) \in B_2\}$ is also an a -simulation. This is sufficient to demonstrate transitivity. The Church-Rosser property states that if $N_t \succeq_a M_s$ and $N_t \succeq_a M'_{s'}$, then there is some model $N'_{t'}$ such that $M_s \succeq_a N'_{t'}$ and $M'_{s'} \succeq_a N'_{t'}$. From Definition 2.2 it follows that M_s and $M'_{s'}$ must be bisimilar to one another with respect to $A - \{a\}$. We may therefore construct such a model $N'_{t'}$ by taking M_s (or $M'_{s'}$) and setting $R_a^{N'} = \emptyset$ and $R_b^{N'} = R_b^M$ for all $b \in A - \{a\}$. It can be seen that $N'_{t'} = (S^M, R^{N'}, V^M, s)$ satisfies the required properties. \square

Finally, note that if $N_t \succeq_a M_s$ and $M_s \succeq_a N_t$ it is not necessarily the case that $M_s \leftrightarrow N_t$.

For example, consider the one agent models M and N where:

- $S^M = \{1, 2, 3\}$, $R_a^M = \{(1, 2), (2, 3)\}$ and $V^M(p) = \emptyset$ for all $p \in P$; and
- $S^N = \{a, b, c, d\}$, $R_a^N = \{(a, b), (b, c), (a, d)\}$ and $V^M(p) = \emptyset$ for all $p \in P$.

These two models are clearly not bisimilar, although $M_1 \preceq N_a$ via $\{(1, a), (2, b), (3, c)\}$ and $N_a \preceq M_1$ via $\{(a, 1), (b, 2), (c, 3), (d, 2)\}$.

3 Syntax and semantics

Assuming an interpretation where different \square_a operators stand for different epistemic operators (each describing what an agent knows), *future event logic* is able express what informative events are consistent with a given information state. The syntax and the semantics of future event logic are as follows.

Definition 3.1 [Language of $\mathcal{L}_{\triangleright}$] Given a finite set of agents A and a set of propositional atoms P , the language of $\mathcal{L}_{\triangleright}$ is inductively defined as

$$\phi ::= p \mid \neg\phi \mid (\phi \wedge \phi) \mid \square_a\phi \mid \blacktriangleright_a\phi$$

where $a \in A$ and $p \in P$.

Standard abbreviations include: $\phi \vee \psi$ iff $\neg(\neg\phi \wedge \neg\psi)$; $\phi \rightarrow \psi$ iff $\neg\phi \vee \psi$; $\diamond_a\phi$ iff $\neg\square_a\neg\phi$. We write $\triangleright_a\phi$ for $\blacktriangleright_a\neg\phi$. We propose a dynamic modal way to interpret the refinement quantification. This means that our future is the *computable future*: $\triangleright_a\phi$ is true now, iff there is an (unspecified) informative event for agent a , or a -refinement, after which ϕ is true.

Definition 3.2 [Semantics of future event logic] Assume an epistemic model $M =$

(S, R, V) . The interpretation of $\phi \in \mathcal{L}_\triangleright$ is defined by induction.

$$\begin{aligned} M_s &\models p \text{ iff } s \in V_p \\ M_s &\models \neg\phi \text{ iff } M_s \not\models \phi \\ M_s &\models \phi \wedge \psi \text{ iff } M_s \models \phi \text{ and } M_s \models \psi \\ M_s &\models \Box_a\phi \text{ iff for all } t \in S : (s, t) \in R_a \text{ implies } M_t \models \phi \\ M_s &\models \blacktriangleright_a\phi \text{ iff for all } M'_{s'} : M_s \succeq_a M'_{s'} \text{ implies } M'_{s'} \models \phi \end{aligned}$$

The logic without the refinement quantifier \blacktriangleright_a is the logic \mathcal{L} of multi-agent epistemic logic.

In other words, $\blacktriangleright_a\phi$ is true in an epistemic state iff ϕ is true in all of its *a-refinements*. Note the inverse direction in the definition: the future epistemic state refines the current epistemic state. Typical model operations that produce an *a-refinement* are: blowing up the model (to a bisimilar model) such as adding copies that are indistinguishable from the current model and one another, removing states accessible only by agent *a*, and removing pairs of the accessibility relation for the agent *a*. Validity in a model, and validity, are defined as usual. For an extended discussion of these semantics and a comparison to related logics see [21].

Lemma 3.3 *The logic $\mathcal{L}_\triangleright$ is bisimulation invariant.*

Proof. This is straightforward, noting \Box_a is bisimulation invariant, and the new operator \blacktriangleright_a is clearly bisimulation invariant since *a-simulation* is transitive and bisimulation is just a specific type of simulation. Therefore, if $M_s \leftrightarrow N_t$, and O_u is any model such that $O_u \preceq_a M_s$ then $O_u \preceq_a N_t$, so by Lemma 2.3, we have $O_u \preceq_a N_t$. Thus, $N_t \models \blacktriangleright_a\phi$ implies $M_s \models \blacktriangleright_a\phi$. The reverse direction is symmetric. \square

Additionally, we may define $\mathcal{L}_\triangleright^\mu$, by including the fixed-point operators μ and ν . Specifically:

Definition 3.4 [Language of $\mathcal{L}_\triangleright^\mu$] Given a finite set of agents A and a set of propositional atoms P , the language of $\mathcal{L}_\triangleright^\mu$ is inductively defined as

$$\phi ::= p \mid \neg\phi \mid (\phi \wedge \psi) \mid \Box_a\phi \mid \blacktriangleright_a\phi \mid \mu x.\phi$$

where $a \in A$, $p \in P$, and the atom x only occurs positively (i.e. in the scope of an even number of negations) in the formula ϕ . We will refer to such an atom x as a *fixed-point variable*. The formula $\nu x.\phi$ is an abbreviation for $\neg\mu x.\neg\phi[\neg x \setminus x]$.

The restriction of this logic to the fragment without refinement quantifiers (the modal μ -calculus) will be referred to as \mathcal{L}^μ . An important technical definition we require is that of a *disjunctive formula*. Let Γ be a finite set of $\mathcal{L}_\triangleright^\mu$ formulas. We let the *cover* operator $\nabla_a\Gamma$ be an abbreviation for $\Box_a \bigvee_{\gamma \in \Gamma} \gamma \wedge \bigwedge_{\gamma \in \Gamma} \Diamond_a \gamma$. (To avoid ambiguity, we note $\bigvee_{\gamma \in \emptyset} \gamma$

is always false, whilst $\bigwedge_{\gamma \in \emptyset} \gamma$ is always true). This operator has previously been used in the definition of disjunctive formulae [8], and has recently been axiomatized [5]. We also note its dual may be written $\Delta_a \Gamma$ as an abbreviation for $\diamond_a \bigwedge_{\gamma \in \Gamma} \gamma \vee \bigvee_{\gamma \in \Gamma} \Box_a \gamma$.

Definition 3.5 [Disjunctive formula] A *disjunctive formula (df)* is specified by the following abstract syntax:

$$\alpha ::= x \mid \alpha \vee \alpha \mid \mu x. \alpha \mid \nu x. \alpha \mid \pi \wedge \overline{\nabla \Gamma} \mid \blacktriangleright_a \alpha \mid \triangleright_a \alpha$$

where π is a conjunction of free literals (atoms or negated atoms, but not fixed-point variables), and $\overline{\nabla \Gamma}$ is an abbreviation for $\nabla_{a_1} \Gamma_{a_1} \wedge \dots \wedge \nabla_{a_n} \Gamma_{a_n}$ such that a_1, \dots, a_n are distinct elements of the set A , and each Γ_{a_i} is a finite set of disjunctive formulas. To avoid ambiguity we may refer to the disjunctive formulas of \mathcal{L}^μ (the ones without \blacktriangleright_a or \triangleright_a operators) as μ -disjunctive formulas.

Proposition 3.6 *Every formula ϕ of \mathcal{L}^μ is equivalent to a μ -disjunctive formula,*

This is shown in [13].

Example 3.7 [Knowledge and belief] Given are two agents that are uncertain about the value of a fact p , and where this is common knowledge, and where p is true. We assume that both accessibility relations are equivalence relations, and that the epistemic operators model the agents' knowledge. An informative event is possible after which a knows that p but b does not know that; this is expressed by

$$\triangleright(\Box_a p \wedge \neg \Box_b \Box_a p)$$

In Figure 1, the structure is on the left, and its refinement validating the post-condition is on the right. In this visualization, the actual state is the (bottom) right one, and states that are indistinguishable for an agent are linked and labeled with the name of that agent, and transitivity and reflexivity are assumed (so on the right, all three states are indistinguishable for agent b). Note that on the left, the formula $\triangleright(\Box_a p \wedge \neg \Box_b \Box_a p)$ is true, because $\Box_a p \wedge \neg \Box_b \Box_a p$ is true in the right structure: in the actual state there is no alternative for agent a , so $\Box_a p$ is true, whereas agent b considers it possible that the top-state is the actual state, and in that state agent a considers it possible that p is false. Therefore, $\neg \Box_b \Box_a p$ is also true in the bottom right state.

Example 3.8 [Controller synthesis] Consider a discrete-event system S to be controlled, with two possible actions c and u . Given a control objective ϕ expressed in, say the μ -calculus, the following formulas express respectively the well-known verification/synthesis problems:

- Controller synthesis: Assume action c is controllable as opposed to u . The system S is *controllable for ϕ* if and only if, $S \models \triangleright_c \phi$, as a c -refinement of S denotes the result of applying some control acting on action c .
- Module checking [14]: The system S is interpreted as an *open system* where action c is internal and action u comes from the environment. The system S satisfies ϕ whatever




Fig. 1. The refinement in Example 3.7.

the environment does but blocking if, and only if, $S \models \blacktriangleright_u(\text{NonBlocking} \Rightarrow \phi)$, where $\text{NonBlocking} \equiv \nu x. \diamond_u \top \wedge \square x$ is an invariant telling that there always exists an environment reaction. The u -refinements with hypothesis NonBlocking denotes all possible non-blocking environments.

- Advanced controller synthesis: We combine the two cases above. We reconsider the control problem for ϕ where S is interpreted as an open system. The open system S can be controlled to achieve property ϕ if, and only if,

$$S \models \triangleright_c \blacktriangleright_u(\text{NonBlocking} \Rightarrow \phi)$$

Example 3.9 [Program logic] Consider a specification, MUTEX , of a mutual exclusion protocol and some property ϕ of this protocol specified in CTL . Now we may ask if we can find a refinement of MUTEX that satisfies ϕ but also such that if Process i is in the critical section ($cs(i)$) at time $n + 1$, then this is known at time n . This is expressed as

$$\text{MUTEX} \models \triangleright(AG[EXcs(i) \Rightarrow AXcs(i)] \wedge \phi)$$

The refinement consists in moving the nondeterministic choices forward, so that a fork at time n becomes a fork at time $n - 1$ with each branch having a single successor at time n .

We also note that Section 6.2 presents an application of the refinement quantification to two-player asynchronous games.

4 Axiomatization: $\mathcal{L}_\triangleright$

Here we present a series of axioms for the logic $\mathcal{L}_\triangleright$. We will derive a number of validities, show the axioms to be sound, and discuss a general strategy for showing their completeness. For simplicity, we will present the axiomatization in the single agent case (and hence the \square_a operator will simply be referred to as \square), although the axiomatization and proofs easily generalize to the multi-agent case. We will also use the relation R simply as a set of pairs $\subseteq S^M \times S^M$, and use the abbreviation $sR^M = \{u \in S^M \mid (s, u) \in R^M\}$.




Fig. 2. The refinement in Example 3.9.

The axiomatization presented is a substitution schema, since the substitution rule itself is not valid. Note that for all atomic propositions p , $p \rightarrow \blacktriangleright p$, but the same is not true for an arbitrary formula (substitute $\diamond\top$ for p in the formula of Example 3.7). This propositional case itself is presented as axiom **G1** and prevents the logic **FEL** from being a normal modal logic.

Definition 4.1 The axiomatization **FEL** is such that the axioms are all substitution instances of the following:

- P** All tautologies of propositional logic
- K** $\Box(\phi \rightarrow \psi) \rightarrow \Box\phi \rightarrow \Box\psi$
- G0** $\blacktriangleright(\phi \rightarrow \psi) \rightarrow \blacktriangleright\phi \rightarrow \blacktriangleright\psi$
- G1** $\alpha \leftrightarrow \blacktriangleright\alpha$ where α is a propositional formula
- GK** $\bigwedge_{\gamma \in \Gamma} \diamond\triangleright\gamma \leftrightarrow \triangleright\bigvee\Gamma$

along with the rules:

- MP** From $\vdash \phi \rightarrow \psi$ and $\vdash \phi$ infer $\vdash \psi$
- Nec1** From $\vdash \phi$ infer $\vdash \Box\phi$
- Nec2** From $\vdash \phi$ infer $\vdash \blacktriangleright\phi$

The axiomatization **K** for the logic \mathcal{L} consists of the axioms **P**, **K**, and the rules **MP** and **Nec1**.

The axiomatization is surprisingly simple given the complexity of the semantic definition of the refinement quantification, \blacktriangleright . We note that while refinement is known

to be reflexive, transitive and satisfies the Church-Rosser property (Lemma 2.3), the corresponding modal axioms are not required. Rather, these properties may be inferred from the axioms presented above.

4.1 Example of derivation

We present a simple derivation of $\diamond\top \rightarrow \triangleright(\nabla\{p\} \vee \nabla\{\neg p\})$. In some cases several deductions have been combined into single statements, but this is restricted to cases of well known modal theorems.

1. **P, Nec1, K** $\vdash \diamond\top \leftrightarrow \diamond(p \vee \neg p)$
2. **P, Nec1, K** $\vdash \diamond(p \vee \neg p) \leftrightarrow (\diamond p \vee \diamond \neg p)$
3. See below $\vdash \diamond p \rightarrow \triangleright \nabla\{p\}$
4. See below $\vdash \diamond \neg p \rightarrow \triangleright \nabla\{\neg p\}$
5. **P, Nec2, G0** $\vdash \triangleright \square p \rightarrow \triangleright(\nabla\{p\} \vee \nabla\{\neg p\})$
6. **P, Nec2, G0** $\vdash \triangleright \square \neg p \rightarrow \triangleright(\nabla\{p\} \vee \nabla\{\neg p\})$
7. **P, MP** $\vdash \diamond\top \rightarrow \triangleright(\nabla\{p\} \vee \nabla\{\neg p\})$

Lines 3 and 4 above require the following deduction, which is true for all propositional formula α :

1. **G1** $\vdash \alpha \leftrightarrow \triangleright \alpha$
2. **P, Nec1, K** $\vdash \diamond \alpha \leftrightarrow \diamond \triangleright \alpha$
3. **GK** $[\Gamma = \{\alpha\}] \vdash \diamond \alpha \leftrightarrow \triangleright \nabla\{\alpha\}$

4.2 Soundness

For notational convenience, given a finite set of $\mathcal{L}_\triangleright$ formulas, $\Gamma = \{\phi_1, \dots, \phi_n\}$, we let $\triangleright\Gamma = \{\triangleright\phi \mid \phi \in \Gamma\}$ (and likewise for other unary operators).

Theorem 4.2 *The axiomatization FEL is sound for $\mathcal{L}_\triangleright$.*

Proof. As all models of $\mathcal{L}_\triangleright$ are models of \mathcal{L} , the schemas **P**, **K** and the rule **MP** and **Nec1** are all sound. We deal with the remaining schemas and rules below:

- G0** Suppose that M_s is a model such that $M_s \models \blacktriangleright(\phi \rightarrow \psi)$. Then for every N_t , where $N_t \preceq M_s$, we have $N_t \models \phi \rightarrow \psi$. Therefore if it is also the case that for every N_t where $N_t \preceq M_s$, we have $N_t \models \phi$, then it follows that every such model also satisfies $\blacktriangleright\psi$.
- G1** Suppose that α is a propositional formula. By Definition 2.2 for every model $N_t \preceq M_s$, for every propositional atom p , we have $s \in V^M(p)$ if and only if $t \in V^N(p)$. As the interpretation of α depends solely on the the valuation of propositions at s , then $M_s \models \alpha$ if and only if $N_t \models \alpha$ for every $N_t \preceq M_s$.
- GK** Suppose M_s is a model such that for some set Γ , $M_s \models \bigwedge_{\gamma \in \Gamma} \diamond \triangleright \gamma$. Therefore for every $\gamma \in \Gamma$ there is some $t^\gamma \in sR^M$ such that $M_{t^\gamma} \models \triangleright \gamma$. Thus, for each $\gamma \in \Gamma$,

there is some model $N_{u^\gamma}^\gamma \preceq M_{t^\gamma}$ such that $N_{u^\gamma}^\gamma \models \gamma$. Without loss of generality, we may assume that for each $\gamma \in \Gamma$ the models N^γ are disjoint. We construct the model M^Γ such that $S^{M^\Gamma} = S^M \cup \bigcup_{\gamma \in \Gamma} S^{N^\gamma}$, $R^{M^\Gamma} = \{(s, u^\gamma) \mid \gamma \in \Gamma\} \cup \bigcup_{\gamma \in \Gamma} R^{N^\gamma}$, and for all $p \in P$, $V^{M^\Gamma}(p) = V^M(p) \cup \bigcup_{\gamma \in \Gamma} V^{N^\Gamma}(p)$.

We can see that $M_s^\Gamma \preceq M_s$, via the relation $\mathcal{R}^\Gamma = \{(s, s)\} \cup \bigcup_{\gamma \in \Gamma} \mathcal{R}^\gamma$ where \mathcal{R}^γ is the refinement relation corresponding to $N_{u^\gamma}^\gamma \preceq M_{t^\gamma}$. Furthermore, for each $t \in sR^{M^\Gamma}$ it is clear that $M_t^\Gamma \Leftrightarrow N_{u^\gamma}^\gamma$ for some γ , and thus $M_t^\Gamma \models \gamma$. Therefore $M_s^\Gamma \models \Box \bigvee_{\gamma \in \Gamma} \gamma$. Finally, for each $\gamma \in \Gamma$ there is some $u^\gamma \in sR^{M^\Gamma}$ where $M_{u^\gamma}^\Gamma \models \gamma$ so $M_s^\Gamma \models \bigwedge_{\gamma \in \Gamma} \Diamond \gamma$. Therefore $M_s^\Gamma \models \nabla \Gamma$, so $M_s \models \triangleright \nabla \Gamma$.

Conversely, suppose that $M_s \models \triangleright \nabla \Gamma$. Therefore, there is a model, $N_t \preceq M_s$ such that $N_t \models \nabla \Gamma$. Expanding the definitions, we have, for every $\gamma \in \Gamma$ there is some $u \in tR^N$ such that $N_u \models \gamma$, and for every $u \in tR^N$ there is some $v \in sR^M$ such that $N_u \preceq M_v$. Combining these statements we have, for every $\gamma \in \Gamma$ there is some $v \in sR^M$ such that $M_v \models \triangleright \gamma$, and thus $M_s \models \bigwedge_{\gamma \in \Gamma} \Diamond \triangleright \gamma$.

Nec2 If ϕ is a validity, then it is satisfied by every model, so for any model M_s , ϕ is satisfied by every model $N_t \preceq M_s$, and hence every model satisfies $\blacktriangleright \phi$. □

4.3 Completeness

We first show that every $\mathcal{L}_\triangleright$ formula is logically equivalent to a \mathcal{L} formula. We then show that if the latter is a theorem in **K**, the former is also a theorem, in **FEL**.

Lemma 4.3 *Every formula of $\mathcal{L}_\triangleright$ is logically equivalent to a formula of \mathcal{L} .*

Proof. As the axiom **GK** is formulated in terms of the cover operator, it is convenient to prove this equivalence by means of an equally expressive version of the modal logic \mathcal{L} that is also formulated with the cover operator [5].² (A direct proof in our own setting is quite possible, but considerably longer.) Consider the syntax of cover logic

$$\phi ::= \perp \mid \top \mid \phi \vee \phi \mid p \wedge \phi \mid \neg p \wedge \phi \mid \nabla \Gamma.$$

The semantics of $\nabla \Gamma$ is the obvious one if we recall our introduction by abbreviation of the cover operator: $M_s \models \nabla \Gamma$ iff for all $\phi \in \Gamma$ there is a $t \in R(s)$ such that $M_t \models \phi$, and for all $t \in R(s)$ there is a $\phi \in \Gamma$ such that $M_t \models \phi$. The modal box and diamond are definable as: $\Box \phi$ iff $\nabla \emptyset \vee \nabla \{\phi\}$, and $\Diamond \phi$ iff $\nabla \{\phi, \top\}$

Now consider the extension of cover logic with the refinement quantification \triangleright . By the definition of \Diamond in cover logic, axiom **GK** now takes shape $\bigwedge_{\gamma \in \Gamma} \nabla \{\triangleright \gamma, \top\} \leftrightarrow \triangleright \nabla \Gamma$. (And this is clearly also sound.) Given a formula ψ in cover logic with refinement, we prove by induction on the number of the occurrences of \triangleright in ψ that it is equivalent to an \triangleright -free formula, and therefore to a formula in the modal logic \mathcal{L} . The base is trivial. Now assume ψ contains $n + 1$ \triangleright -operators. Choose a subformula of type $\triangleright \phi$ of our given formula ψ , where ϕ is \triangleright -free (i.e. choose an innermost \triangleright). We prove by induction on the structure of ϕ that $\triangleright \phi$ is logically equivalent to a formula χ without \triangleright .

² We thank Yde Venema for suggesting this proof.

- $\triangleright \perp$ iff \perp .
- $\triangleright \top$ iff \top .
- $\triangleright(p \wedge \phi)$ iff $p \wedge \triangleright\phi$ (refinements do not affect atoms); **IH**.
- $\triangleright(\neg p \wedge \phi)$ iff $\neg p \wedge \triangleright\phi$ (refinements do not affect atoms); **IH**.
- $\triangleright(\phi \vee \psi)$ iff $\triangleright\phi \vee \triangleright\psi$ (directly from the semantics of \triangleright); **IH**.
- $\triangleright\nabla\Gamma$ iff $\bigwedge_{\gamma \in \Gamma} \nabla\{\triangleright\gamma, \top\}$; **IH**. (By induction, each $\triangleright\gamma$ is equivalent to an \triangleright -free formula ψ , and the resulting $\bigwedge_{\psi} \nabla\{\psi, \top\}$ is also \triangleright -free.)

Thus we are able to push the refinement operators deeper into the formula until they eventually reach \top or \perp , at which point they disappear and we are left with χ (which does not contain \triangleright and is equivalent to $\triangleright\phi$). Replacing $\triangleright\phi$ by χ in ψ gives a result with at least one less \triangleright -operator, to which the (original) induction hypothesis applies. \square

Lemma 4.4 *Let $\phi \in \mathcal{L}_{\triangleright}$ be given and $\psi \in \mathcal{L}$ be equivalent to ϕ . If ψ is a theorem in **K**, then ϕ is a theorem in **FEL**.*

Proof. Given a $\phi \in \mathcal{L}_{\triangleright}$, Lemma 4.3 gives us an equivalent $\psi \in \mathcal{L}$. Assume that ψ is a theorem in **K**. We can extend the derivation of ψ to a derivation of ϕ by observing that the first five of the six itemized reduction steps in Lemma 4.3 are all provable equivalences, and that the last item is of course the axiom **GK**. (Where we also need to observe that the system **FEL** satisfies the substitution of equivalents: if ϕ_1 is equivalent to ϕ_2 and ϕ_1 is a subformula of ϕ_3 , and ϕ_3 is a theorem, then $\phi_3[\phi_1 \setminus \phi_2]$ is also a theorem.) \square

Theorem 4.5 *The axiom schema **FEL** is sound and complete for the logic $\mathcal{L}_{\triangleright}$.*

Proof. The soundness proof is given in Theorem 4.2, so we are left to show completeness. Suppose that ϕ is valid: $\models \phi$. Applying Lemma 4.3 we know that there is some equivalent formula ψ not containing any refinement quantification. As ϕ is valid, from that and the validity $\phi \leftrightarrow \psi$ it follows that ψ is also valid in future event logic, and therefore also valid in the logic \mathcal{L} (note that the model class is the same!). From the completeness of **K** it follows that ψ is derivable, i.e. it is a theorem. From Lemma 4.4 it follows that ϕ is a theorem. \square

5 Axiomatization: $\mathcal{L}_{\triangleright}^{\mu}$

The axiomatization for $\mathcal{L}_{\triangleright}^{\mu}$ extends the axiomatization for $\mathcal{L}_{\triangleright}$ with the extra axiom and rule of Kozen's axiomatization of the modal μ -calculus (**F1** and **F2**), and two new interaction axioms (**G3** and **G4**). The axiomatization **FEL** _{μ} is a substitution schema of the axioms and rules for $\mathcal{L}_{\triangleright}$, **FEL** (see Section 4), along with the axiom and rule for the modal μ -calculus:

$$\mathbf{F1} \quad \phi[\mu x.\phi \setminus x] \rightarrow \mu x.\phi$$

$$\mathbf{F2} \quad \text{From } \phi[\psi \setminus x] \rightarrow \psi \text{ infer } \mu x\phi \rightarrow \psi$$

and two new interaction axioms:

$$\mathbf{G3} \triangleright \mu x.\phi \leftrightarrow \mu x.\blacktriangleright\phi \text{ where } \mu x.\phi \text{ is a } df \text{ (Def. 3.5)}$$

$$\mathbf{G4} \triangleright \nu x.\phi \leftrightarrow \nu x.\blacktriangleright\phi \text{ where } \nu x.\phi \text{ is a } df$$

These interaction axioms have an important associated condition: the refinement quantification will only commute with a fixed-point operator if the fixed-point formula is a disjunctive formula.

5.1 Soundness

The soundness proofs of Section 4.2 still apply and the soundness of **F1** and **F2** are well known [3], so we are left to show that **G3** and **G4** are sound.

Theorem 5.1 *The axioms G3 and G4 are sound.*

Proof. In this proof we will find it convenient to use the bisimulation quantifiers characterization of both fixed-point operators and refinement quantification. We recall what bisimulation quantifier is: Given an atom x and a formula ϕ , the expression $\exists x\phi$ means that there exists x such that ϕ , and it is interpreted as $M_s \models \exists x\phi$ if, and only if, for some N_t bisimilar to M_s *except for* x —for which we will write $N_t \Leftrightarrow_x M_s$ — we have $N_t \models \phi$. We let $\forall x\phi$ abbreviate $\neg\exists x\neg\phi$, and a deeper technical discussion of the properties of bisimulation quantifiers may be found in [8].

- (i) $\mu x.\phi$ is equivalent to $\forall x(\blacksquare(\phi \rightarrow x) \rightarrow x)$ [11] (where \blacksquare is the universal modality which quantifies over all states in the model).
- (ii) $\nu x.\phi$ is equivalent to $\exists x(\blacksquare(x \rightarrow \phi) \wedge x)$ [11].
- (iii) $\blacktriangleright\phi$ is equivalent to $\forall r\phi^r$, where ϕ^r is the relativization of ϕ to the atom r , which may be computed recursively by replacing every occurrence of $\square\psi$ in ϕ with $\square(r \rightarrow \psi^r)$ [21].
- (iv) $\triangleright\phi$ is equivalent to $\exists r\phi^r$.

Note that from [7] we know that bisimulation quantifiers are expressible in the modal μ -calculus, and thus the equivalences (i) and (ii) hold in the modal μ -calculus. Furthermore, in [21], the equivalences (iii) and (iv) are shown to hold for all logics that are closed under bisimulation and announcement. As the modal μ -calculus is such a logic, all four equivalences hold in the modal μ -calculus, and they may be reasonably applied in the proofs given below:

G3 It is more convenient in this proof to reason about the axiom in its contrapositive form: $\triangleright \nu x. \phi \leftrightarrow \nu x. \triangleright \phi$. Using the equivalent transformations above we have:

$$\begin{aligned}
\triangleright \nu x. \phi &\leftrightarrow \exists r \exists x (\blacksquare(x \rightarrow \phi) \wedge x)^r \\
&\leftrightarrow \exists x \exists r (\blacksquare(x \rightarrow \phi^r) \wedge x) \\
&\leftrightarrow \exists x (\exists r \blacksquare(x \rightarrow \phi^r) \wedge x) \\
&\rightarrow \exists x (\blacksquare \exists r (x \rightarrow \phi^r) \wedge x) \\
&\leftrightarrow \exists x (\blacksquare(x \rightarrow \exists r \phi^r) \wedge x) \\
&\leftrightarrow \nu x. \triangleright \phi
\end{aligned}$$

This proof simply applies known validities of bisimulation quantifiers. Note that the fourth line is not an equivalence in the general case. However, we may show that where ϕ is a disjunctive formula, the equivalence does hold. To do this, suppose M_s is any countable model such that $M_s \models \exists x (\blacksquare \exists r (x \rightarrow \alpha^r) \wedge x)$, where α is a disjunctive formula. As the μ -calculus enjoys the tree-model property, we may assume that there is some tree-like model $N_t \stackrel{i}{\leftrightarrow}_x M_s$ such that $N_t \models x \wedge \blacksquare \exists r (x \rightarrow \alpha^r)$. We inductively build a series of models $N_t^i \stackrel{i}{\leftrightarrow}_{r,x} N_t$ where $N^i = (S^N, R^N, V_i)$. We set $V_0(x) = \{t\}$, $V_0(r) = \emptyset$ and $V_0(p) = V^N(p)$ for all $p \notin \{r, x\}$. As $N_t \models \exists r \alpha^r$ and $\nu x. \alpha$ is a disjunctive formula, the only case where the atom x may influence the interpretation of $\exists r \alpha^r$ is at a set of states such that all states beyond that set of states are irrelevant to the interpretation of $\exists r \alpha^r$ at t (this set of states forms a *frontier*). This is because from Definition 3.5, if x is a sub-formula of α , then if x appears in the scope of a conjunction, it appears within the scope of a modality within that conjunction. Thus, there is a set of states $\{u_0, u_1, \dots\} \in V^N(x)$ such that $N_t^i \models \alpha^r$, where $N^i = (S^N, R^i, V^i)$ for $V^i(x) = \{t, u_0, u_1, \dots\}$, $V^i(y) = V^N(y)$ for $y \notin \{x, r\}$ and $R^i = R^N \setminus \{(u_i, s) \mid s \in S^N, i = 0, 1, \dots\}$. Consequently the valuation of r maybe restricted to states that are not reachable from any state, $\{u_0, u_1, \dots\}$. Let $S_0 \subset S^N$ be the set of states reachable from t , but not reachable from u_i for any i . We define N^1 by setting $V_1(x) = V^1(x)$, $V_1(r) = V^1(r) \cap S_0$ and $V_1(y) = V^N(y)$ for $y \notin \{x, r\}$. As $u_0, u_1, \dots \in V^N(x)$, we have $N_{u_i} \models x \wedge \blacksquare (\exists r (x \rightarrow \alpha^r))$ for all i .

As M_s is a countable model, we may assume an enumeration of the worlds (or states) in that model. The induction proceeds by taking the first state u_0 on the frontier and repeating the process (i.e. finding a valuation V' such that V' make x true on a frontier $\{v_0, v_1, \dots\}$, agrees with V^N on the interpretation of all atoms except x and r , makes $N_{u_0}^i \models \alpha^r$ and makes $N_{v_i} \models x \wedge \blacksquare (\exists r (x \rightarrow \alpha^r))$ for all i). We define V_2 by taking the union of $V_2(x) = V^1(x)$ and $V_2(r) = V_1(r) \cup (V^1(r) \cap S_1)$ where S_1 is the set of states reachable from u_0 , but not from v_i for and i , and all other atoms have their valuations unchanged. The states $\{v_0, v_1, \dots\}$ are added to the set of frontier states and the induction continues. As the sets $V_i(x)$ and $V_i(r)$ are strictly increasing with i , this process is well defined, and its limit N_t^* will satisfy $\exists x \exists r (\blacksquare(x \rightarrow \alpha^r) \wedge x)$, as required. The construction is represented in Figure 3.

G4 We also use the contrapositive form of the axiom: $\triangleright \mu x. \phi \leftrightarrow \mu x. \triangleright \phi$.




Fig. 3. The inductive step for the construction of N_t^* . The formula α^r is independent of any state where r is not true, or any state beyond the frontier defined by u_0, u_1, \dots

$$\begin{aligned}
\triangleright \mu x. \phi &\leftrightarrow \exists r (\forall x (\blacksquare (\phi \rightarrow x) \rightarrow x))^r \\
&\rightarrow \forall x \exists r (\blacksquare (\phi \rightarrow x) \rightarrow x)^r \\
&\leftrightarrow \forall x \exists r (\blacklozenge (\phi \wedge \neg x) \vee x)^r \\
&\leftrightarrow \forall x (\exists r \blacklozenge (\phi \wedge \neg x)^r \vee x) \\
&\leftrightarrow \forall x (\blacklozenge \exists r (\phi \wedge \neg x)^r \vee x) \\
&\leftrightarrow \forall x (\blacklozenge (\exists r \phi^r \wedge \neg x) \vee x) \\
&\leftrightarrow \forall x (\blacksquare (\exists r \phi^r \rightarrow x) \rightarrow x) \\
&\leftrightarrow \mu x. \triangleright \phi
\end{aligned}$$

While the forward implication is generally true, the right to left implication again relies on the fact that $\mu x \phi$ is a disjunctive formula. For this, write the formula $\mu x \alpha$ to emphasize it is a disjunctive formula. We use the inductive characterization of $\mu x. \triangleright \alpha$: if for $x \in P$ and $S \subseteq S^M$, $M^{[x \mapsto S]} = (S^M, R^M, V)$ such that $V(x) = S$ and $V(y) = V^M(y)$ for $y \neq x$, then we may inductively define $\|\triangleright \alpha\|_0 = \emptyset$, and $\|F\alpha\|_i = \{s \in S^M \mid M_s^{[x \mapsto \bigcup_{j < i} \|\triangleright \alpha\|_j]} \models \triangleright \alpha\}$. Then $M_s \models \mu x. \triangleright \alpha$ if and only if $s \in \|\triangleright \alpha\|_\tau$, where τ is an ordinal [3].

Now suppose $M_s \models \mu x. \triangleright \alpha$. Without loss of generality we may suppose that M is a countable tree-like model. As M_s satisfies $\mu x. \triangleright \alpha$, there must be some least ordinal τ whereby $s \in \|\triangleright \alpha\|_\tau$. We give a proof by induction, and the base case where $\tau = 0$ is trivial. Let $M^\tau = M^{[x \mapsto \bigcup_{j < \tau} \|\triangleright \alpha\|_j]}$, and then $M_s^\tau \models \triangleright \alpha$. As $\mu x \alpha$ is a disjunctive formula, we are again in the case where there is a refinement of M^τ and a frontier such that x may only be true at s or on the frontier, and no point beyond the frontier affects the interpretation of α . Formally, there is a set of states $\{u_0, u_1, \dots\} \in V^{M^\tau}(x)$ such that $M'_s \models \exists r \alpha^r$, where $M' = (S', R', V')$ such that

- $S' \subseteq S^{M^\tau}$ is the set of states reachable from s , but not from any u_i ;
- $V'(x) = \{t, u_0, u_1, \dots\}$, $V'(y) = V^{M^\tau}(y)$ for $y \neq x$; and

- $R' = R^{M^\tau} \setminus \{(u_i, t) \mid t \in S^{M^\tau}, i = 0, 1, \dots\}$.

We note that M'_s is a refinement of M_s^τ . Now as for each i , $u_i \in \|\triangleright\alpha\|_j$ for some $j < \tau$, by the inductive hypothesis we may assume there is some model $N^i = (S^i, R^i, V^i)$ where $N_{v_i}^i \preceq M_{u_i}^\tau$ and $N_{u_i}^i \models \mu x.\alpha$. We may append these models to M' to define $M^* = (S^*, R^*, V^*)$ where $S^* = S' \cup \bigcup_i S^i$, $R^* = R' \cup \bigcup_i R^i \cup \{(t, v_i) \mid (t, u_i) \in R'\}$, and $V^*(y) = V'(y) \cup \bigcup_i V^i(y)$ for all $y \in P$. It is clear that M_s^* is a refinement of M_s , and by the axiom **F1** we can see $M_s^* \models \mu x.\alpha$ as required. \square

We note that the general form of **G3** is not sound, for example, take $\phi = \mu z.\diamond(y \rightarrow z) \rightarrow \diamond(\neg y \rightarrow x)$. Then $\blacktriangleright\mu x.\phi$ is true if y is true at every immediate successor of the current state, whereas $\mu x.\blacktriangleright\phi$ is only true at states with no successor. Likewise **G4** is not true in the general case, as can be seen by taking $\phi = p \wedge \square(\diamond\top \rightarrow x)$. Then $\nu x.\blacktriangleright\phi$ is true if and only if p is true at every reachable state, and $\blacktriangleright\nu x.\phi$ is true only if p is true at every state within one step.

5.2 Completeness

The completeness proof of \mathbf{FEL}_μ proceeds exactly as for Theorem 4.5, replacing the formulas in cover logic with disjunctive formulas to get a statement similar to the one of Lemma 4.3.

Lemma 5.2 *Every formula of $\mathcal{L}_\triangleright^\mu$ is equivalent in \mathbf{FEL}_μ to a formula of the modal μ -calculus \mathcal{L}^μ .*

Proof. Given a formula ψ , we prove by induction on the number of the occurrences of \triangleright in ψ that it is equivalent to an \triangleright -free formula, and therefore to a formula in the modal μ -calculus \mathcal{L}^μ . The base is trivial. Now assume ψ contains $n + 1$ \triangleright -operators. Choose a subformula of type $\triangleright\phi$ of our given formula ψ , where ϕ is \triangleright -free (i.e. choose an innermost \triangleright). As ϕ is \triangleright -free, it follows from Proposition 3.6 that ϕ is semantically equivalent to a formula in disjunctive normal form, and by the completeness of Kozen's axiom system [23] this equivalence is provable in \mathbf{FEL}_μ . By **Nec2** and **G0** it follows that $\triangleright\phi$ is provably equivalent to some formula $\triangleright\psi$ where ψ is a disjunctive formula. Thus without loss of generalization, we may assume in the following that ϕ is in disjunctive normal form. We may now proceed by induction over the complexity of ϕ , and conclude that $\triangleright\phi$ is logically equivalent to a formula χ without \triangleright . All cases of this induction are as before, we only show the final two, different cases:

- $\triangleright\mu x.\phi$ iff $\mu x.\triangleright\phi$ (by **G4** noting that all subformulas of a disjunctive formula are themselves disjunctive); IH.
- $\triangleright\nu x.\phi$ iff $\nu x.\triangleright\phi$ (by **G3**); IH.

Replacing $\triangleright\phi$ by χ in ψ gives a result with one less \triangleright -operator, to which the (original) induction hypothesis applies. \square

Theorem 5.3 *The axiom schema \mathbf{FEL}_μ is sound and complete for the logic $\mathcal{L}_\triangleright^\mu$*

Proof. Soundness follows from Theorem 5.1 and Theorem 4.2. To see \mathbf{FEL}_μ is complete, suppose ϕ is a valid formula. Then by Lemma 5.2, ϕ is provably equivalent to some valid formula ψ of \mathcal{L}^μ . As ψ is valid, it must be provable since \mathbf{P} , \mathbf{K} , $\mathbf{F1}$, $\mathbf{F2}$, $\mathbf{Nec1}$, and \mathbf{MP} give a sound and complete proof system for the modal μ -calculus [23]. A proof of ϕ follows by \mathbf{MP} . \square

6 Complexity

Both axiomatizations demonstrated the expressivity and decidability of $\mathcal{L}_\triangleright$ (expressively equivalent to \mathbf{K}) and $\mathcal{L}_\triangleright^\mu$ (expressively equivalent to \mathcal{L}^μ). Decidability for both follows from the fact that a computable translation is given in the completeness proofs. Note that as the translations given are recursive and involve translating formulas to disjunctive normal form, the translation is non-elementary in the size of the original formula. In this section we examine the complexity of $\mathcal{L}_\triangleright$, providing both an elementary upper bound and a succinctness proof.

6.1 Upper-Bound

A decision procedure for $\mathcal{L}_\triangleright$ is given via a tableau. Given any $\mathcal{L}_\triangleright$ formula ϕ , we describe a tableau construction that either constructs a model for ϕ , or reports that ϕ is not satisfiable.

Definition 6.1 A formula is in positive normal form if it is built from the following abstract syntax.

$$\alpha ::= \top \mid \perp \mid p \mid \neg p \mid \alpha \wedge \alpha \mid \alpha \vee \alpha \mid \Box \alpha \mid \Diamond \alpha \mid \blacktriangleright \alpha \mid \triangleright \alpha$$

We note every $\mathcal{L}_\triangleright$ formula may be converted into positive normal form with linear change to the size of formula.

Tableau Definition:

Let ϕ be a formula in positive normal form. Suppose that each subformula of ϕ is uniquely indexed as ϕ_i for $i = 0, \dots, m$ (thus, two identical subformulas appearing in different places in ϕ would be indexed differently). Let $I = \{1, \dots, m\}$, let \subset be the subformula relation over these nodes (so $j \subset i$ if and only if ϕ_j is a subformula of ϕ_i), and given $\sigma \subseteq I$, let σ^+ be the set $\{j \in I \mid \exists i \in \sigma, i \subset j\}$. Suppose also that $\phi = \phi_0$. The initial tableau, $T_0 \in \wp(\wp(I))$ consists of the set of nodes, σ , each of which is a subset of I satisfying the following conditions:

- if $i \in \sigma$ and $\phi_i = \phi_j \wedge \phi_k$ then $j, k \in \sigma$;
- if $i \in \sigma$ and $\phi_i = \phi_j \vee \phi_k$ then either $j \in \sigma$ or $k \in \sigma$;
- if $i \in \sigma$ and $\phi_i = \triangleright \phi_j$ or $\phi_i = \blacktriangleright \phi_j$ then $j \in \sigma$;
- if $i, j \in \sigma$ then if $\phi_i = p$, then $\phi_j \neq \neg p$.

The tableau, T_n is then successively pruned according to a game for each node:

Definition 6.2 Let the two players be \mathfrak{E} and \mathfrak{A} , and σ be some node in T_n . We define the *pruning game* $\mathcal{G}(T_n, \sigma)$ where each game position is a tuple (Θ, i) where $\Theta \subset T_n$ and $i \in I$. For two sets $\Theta_1, \Theta_2 \subset T_n$ we define $\Theta_1 \sqsubseteq \Theta_2$ if and only if for every $\theta \in \Theta_1$ there is some $\Theta' \subset \Theta_1$ and some $\lambda \in \Theta_2$ where $\theta \cup \bigcup_{\rho \in \Theta'} \rho = \lambda$.

Init Player \mathfrak{E} selects some $\Theta \subseteq T_n$, and then the initial state is $(\Theta, 0)$.

Move given the state (Θ, i) :

- (i) if $\phi_i = \blacktriangleright \alpha_j$, \mathfrak{A} selects some $\Theta' \sqsubseteq \Theta$, and the new game position is (Θ', j) ,
- (ii) else if $\phi_i = \triangleright \alpha_j$, \mathfrak{E} selects some $\Theta' \sqsubseteq \Theta$, and the new game position is (Θ', j) ,
- (iii) else if $\phi_i = \phi_j \wedge \phi_k$, \mathfrak{A} selects $\ell \in \{j, k\}$ and the new game position is (Θ, ℓ) ,
- (iv) else if $\phi_i = \phi_j \vee \phi_k$, \mathfrak{E} selects $\ell \in \{j, k\} \cap \sigma^+$ and the new state is (Θ, ℓ) .
- (v) else if $\phi_i \notin \sigma$ and $\phi_i = \square \phi_j$ or $\phi_i = \diamond \phi_j$, the new state is (Θ, j) .

Wins The game proceeds until no further move can be made. For such a game position (Θ, i) :

- (i) if $\phi_i = p, \neg p, \top$ or $i \notin \sigma$, player \mathfrak{E} wins,
- (ii) else if $\phi_i = \perp$, player \mathfrak{A} wins,
- (iii) else if $\phi_i = \square \phi_j$, then if for all $\theta \in \Theta, j \in \theta$, then \mathfrak{E} wins and otherwise \mathfrak{A} wins,
- (iv) else if $\phi_i = \diamond \phi_j$, then if for some $\theta \in \Theta, j \in \theta$, then \mathfrak{E} wins and otherwise \mathfrak{A} wins.

The next tableau is then $T_{n+1} = \{\sigma \in T_n \mid \mathfrak{E} \text{ has a winning strategy in } \mathcal{G}(T_n, \sigma)\}$. We note that each game is easily determined since because the subformula ϕ_i is strictly decreasing, there is a maximum of m moves in any game. Furthermore, T_0 is finite and $T_{n+1} \subseteq T_n$, so a fixed point T^* is eventually reached. If for some $\sigma \in T^*$ we have $0 \in \sigma$ the tableau reports that ϕ is satisfiable, and otherwise it reports ϕ is unsatisfiable. We let $\mathcal{G}(\sigma)$ abbreviate $\mathcal{G}(T^*, \sigma)$.

The intuition behind this tableau is that each node represents a state in the model, and records which parts of the formula ϕ are satisfied at that state. The semantics of the \triangleright and \blacktriangleright operator are captured by a game that is played at each state in the model (i.e. successors may be kept, pruned or split). To take a global view, the players are playing a game over T_n where \mathfrak{E} is trying to show a model for ϕ exists, and \mathfrak{A} is trying to show that whichever model \mathfrak{E} builds does not satisfy ϕ . Every time they get to a new state (i.e. they reach a game position (Θ, i) where $i = \square \phi_j$ or $\diamond \phi_j$) they replay the series of moves that brought them to that state, so that each player may select, in turn, refinements of the set of successors for the new state.

Lemma 6.3 *If the tableau reports that ϕ is satisfiable, then ϕ has a model.*

Proof. Suppose that T^* is the final tableau, and $\sigma \in T^*$ and $0 \in \sigma$. We build a model $M = (S, R, V)$ from T^* where $S = T^*$, for all $\theta \in S$, if $\theta \in V(p)$ if and only if $p \in \theta$, and for each $\theta \in S$, $\{\xi \mid (\theta, \xi) \in R\}$ is the first move of player \mathfrak{E} 's winning strategy in the game $\mathcal{G}(\theta)$. By induction over ϕ we may see that $M_\sigma \models \phi$. For our inductive hypothesis we assume if \mathfrak{E} has a winning strategy for the game position $(\{\xi \mid (\theta, \xi) \in R^M\}, i)$, and $i \in \theta$, then $M_\theta \models \phi_i$. Let θM be the set of successors of θ in the model M .

- (i) If $i \in \theta$ where $\phi_i \in \{p, \neg p, \top\}$ then $M_\theta \models \phi_i$.
- (ii) If $i \in \theta$ where $\phi_i = \phi_j \wedge \phi_k$ then \mathfrak{E} must have a winning strategy for the game

- position $(\theta M, i)$ in $\mathcal{G}(\theta)$, so \mathfrak{E} must also have a winning strategy for $(\theta M, j)$ and a winning strategy $(\theta M, k)$, so by the inductive hypothesis we have $M_\theta \models \phi_i$.
- (iii) If $i \in \theta$ where $\phi_i = \phi_j \vee \phi_k$ then \mathfrak{E} must have a winning strategy for the game position $(\theta M, i)$ in $\mathcal{G}(\theta)$, so \mathfrak{E} must also have a winning strategy for $(\theta M, j)$ or a winning strategy for $(\theta M, k)$, so by the inductive hypothesis we have $M_\theta \models \phi_i$.
 - (iv) If $i \in \theta$ where $\phi_i = \Box \phi_j$ then every $\xi \in \theta M$ must have $j \in \xi$. Therefore \mathfrak{E} has a winning strategy from $(\xi M, j)$ in $\mathcal{G}(\xi)$, so every successor of θ satisfies ϕ_j .
 - (v) If $i \in \theta$ where $\phi_i = \Diamond \phi_j$ then some $\xi \in \theta M$ must have $j \in \xi$. Therefore \mathfrak{E} has a winning strategy from $(\xi M, j)$ in $\mathcal{G}(\xi)$, so some successor of θ satisfies ϕ_j .
 - (vi) If $i \in \theta$ where $\phi_i = \blacktriangleright \phi_j$, then \mathfrak{E} has a winning strategy in the game $\mathcal{G}(\theta)$ from the game position $(\theta M, i)$, so for every $A \sqsubseteq \theta M$, player \mathfrak{E} has a winning strategy from the game position (A, j) . Every refinement M'_θ of M_θ may be represented by the restrictions $A^\xi \sqsubseteq \xi M$ for all ξ reachable from θ . As \mathfrak{E} has a winning strategy for all such game positions starting from (A^θ, j) we have $M'_\theta \models \phi_j$ and the result follows.
 - (vii) If $i \in \theta$ where $\phi_i = \triangleright \phi_j$, then \mathfrak{E} has a winning strategy to select restrictions $A^\xi \sqsubseteq \xi M$ for all ξ reachable from θ , in the game $\mathcal{G}(\xi)$ so that she has a winning strategy from the game position (A^ξ, j) . Collecting these restrictions together we are able to define a single refinement M'_θ of M_θ for which \mathfrak{E} has a winning strategy from $(\theta M', j)$, and thus $M'_\theta \models \phi_j$.

By induction it follows that since \mathfrak{E} has a winning strategy for $\mathcal{G}(\sigma)$, $M_\sigma \models \phi$. \square

Lemma 6.4 *If ϕ is satisfiable, then the tableau reports that ϕ is satisfiable.*

Proof. If ϕ is satisfiable, then ϕ has some model, M_s . Seeing as ϕ is equivalent to a formula of \mathcal{L} we may assume that M_s contains no infinite paths [6]. We use M_s to build a set of nodes in the tableau and define a winning strategy for \mathfrak{E} in each, ensuring they are never pruned. The construction of the tableau mirrors the semantics of $\mathcal{L}_\triangleright$. Given the set of all refinements of M modulo bisimulation, \mathcal{M} , and the states in S^M , we index a set of nodes as n_t^N where $t \in S^M$ and $N \in \mathcal{M}$. We ensure $0 \in n_s^M$ and build up the nodes as follows:

- (i) if $i \in n_t^N$ and $i = \phi_j \wedge \phi_k$, then $j, k \in n_t^N$,
- (ii) if $i \in n_t^N$ and $i = \phi_j \vee \phi_k$, then $j \in n_t^N$ if and only if $N_t \models \phi_j$, and $k \in n_t^N$ if and only if $N_t \models \phi_k$,
- (iii) if $i \in n_t^N$ and $i = \blacktriangleright \phi_j$, then for all $N'_t \preceq N_t$, $j \in n_t^{N'}$ ³,
- (iv) if $i \in n_t^N$ and $i = \triangleright \phi_j$, then for all $N'_t \preceq N_t$, $j \in n_t^{N'}$ if and only if $N'_t \models \phi_j$,
- (v) if $i \in n_t^N$ and $i = \Box \phi_j$, then for all u where $(t, u) \in R^N$, $j \in n_u^N$,
- (vi) if $i \in n_t^N$ and $i = \Diamond \phi_j$, then for all u where $(t, u) \in R^N$, $j \in n_u^N$ if and only if $N_u \models \phi_j$.

³ We assume, without loss of generality that every state in every refinement is associated with a single state from S^M

It is clear from the semantics of $\mathcal{L}_\triangleright$ that that for all $i \in n_t^N$, $N_t \models \phi_i$. Let \mathcal{T} be the set of nodes $\{n_t \mid n_t = \bigcup_{N_t \preceq M_t} n_t^N\}$. Now \mathfrak{E} 's moves may be dictated by the model M_s . For the node n_t , the game $\mathcal{G}(n_t)$ simulates the set of formulas that must be true along the path leading to t (whose index is not in n_t in the model M), and the sets of formula true at M_t (whose index is in n_t). Throughout the play, player \mathfrak{E} records a tuple (N, u) of the current refinement and state that is being used to evaluate the formula ϕ_i , where the game position is (Θ, i) . Furthermore, at each step she may ensure that Θ represents the set of nodes $\{n_v \mid (t, v) \in R^N\}$. As \mathfrak{E} is guided by the semantic interpretation of ϕ in M_s , her recorded tuple (N, u) for the game position (Θ, i) is such that $N_u \models \phi_i$. If at the end of play, $u \neq t$, then it will be the case that $i \notin n_t$, so player \mathfrak{E} wins. Otherwise, we will have the game position (Θ, i) and either:

- (i) $\phi_i \in \{\top, p, \neg p\}$ in which case \mathfrak{E} wins (\perp is not an option since $N_t \not\models \perp$); or
- (ii) $\phi_i = \Box \phi_j$ so $N_t \models \Box \phi_j$ and thus for all u where $(t, u) \in R^N$, $N_u \models \phi_j$ so $j \in n_u = \Theta$ and \mathfrak{E} wins; or
- (iii) $\phi_i = \Diamond \phi_j$ so $N_t \models \Diamond \phi_j$ and thus for some u where $(t, u) \in R^N$, $N_u \models \phi_j$ so $j \in n_u = \Theta$ and \mathfrak{E} wins.

Therefore no node $n_t \in \mathcal{T}$ is pruned from the tableau, and as $0 \in n_s$ the tableau reports that ϕ is satisfiable. \square

Corollary 6.5 *The satisfiability problem for $\mathcal{L}_\triangleright$ can be determined in 2EXP time.*

Proof. This follows directly from the tableau description. If ϕ is a formula of size m , then there at most 2^m nodes in the initial tableau. To do the pruning steps we must search all possible strategies for \mathfrak{E} to see if any are winning strategies. As the players moves involve sets of nodes, this takes time 2^{2^m} . For each step we must examine $\mathcal{G}_n(\sigma)$ for every node σ , and as the tableau are strictly decreasing, there are at most 2^m steps. Thus the overall complexity is $2^{O(2^m)}$. \square

It is not yet known whether this complexity bound is optimal. However, below we show that $\mathcal{L}_\triangleright$ is exponentially more succinct than \mathcal{L}^μ , which suggests that the 2EXP bound may be optimal.

6.2 Succinctness

Here we use the refinement quantification to show that $\mathcal{L}_\triangleright$ is able to express the property that two binary trees are n -bisimilar, with a formula of size $O(n^2)$. We will then show that neither \mathbf{K} , nor \mathcal{L}^μ are able to express this property in size less than $2^{O(n)}$.

The basic idea of this construction is to encode a pebble game (or bisimulation game) for showing n -bisimilarity, using the refinement quantification to encode players moves. We restrict our attention to complete binary trees labelled by a single atom, a that marks a prefix closed subtree, and consider the property: “The left subtree marked by a is n -bisimilar to the right subtree marked by a ”. To enforce the binary nature of the tree we suppose that there is an atom ℓ that labels each left successor, and we suppose that r is an abbreviation for $\neg\ell$. We may then refer to the left successor using the modal abbreviations $\langle \ell \rangle \phi$ for $\Diamond(\ell \wedge \phi)$, and likewise for the right successor. Note orientation

(left or right) of a successor does not affect whether two subtrees are bisimilar. They are just used to ensure that the rules of the game are followed.

Our intent is to encode a pebble game played by a *Spoiler* and a *Duplicator*. Each player takes turns at selecting a successor (or moving a pebble) in either subtree. Spoiler goes first, selecting a successor in either subtree (left or right), where a is true, and then Duplicator must select a successor in the other subtree where a is true. If Spoiler is ever unable to move Duplicator wins, and if Duplicator is unable to move, Spoiler wins. If Duplicator has a strategy to survive at least n moves, then the left and right sub-trees must be n -bisimilar [19].

In $\mathcal{L}_\triangleright$ we simulate “selecting a successor” by taking a refinement that leaves only the left or right successor, but otherwise leave the tree intact. To do this we introduce the abbreviation $trunk_m$ to represent a (1-2)-tree where nodes of height less than m have only a left successor, or only a right successor and nodes of height greater than or equal to m , but less than n have a left successor and a right successor:

$$trunk_m = \bigwedge_{i=1}^m (\Box^i(\ell \wedge a) \vee \Box^i(r \wedge a)) \wedge \bigwedge_{i=m}^{n-1} \Box^i(\diamond\ell \wedge \diamond r)$$

We can present the definition for n -bisimilarity recursively, where:

$$\mathcal{B}_i^n = \triangleright \left[\left(\begin{array}{c} \langle \ell \rangle trunk_i \wedge \langle r \rangle trunk_{i-1} \\ \vee \\ \langle r \rangle trunk_i \wedge \langle \ell \rangle trunk_{i-1} \end{array} \right) \longrightarrow \triangleright [\langle \ell \rangle trunk_i \wedge \langle r \rangle trunk_i \wedge \mathcal{B}_{i+1}^n] \right]$$

and $\mathcal{B}_n^n = \langle \ell \rangle trunk_n \wedge \langle r \rangle trunk_n$. Then the property of n -bisimilarity is just equivalent to \mathcal{B}_1^n .

A game scenario is presented in Figure 4.

Lemma 6.6 *Let M_s be a complete binary tree. Then $M_s \models \mathcal{B}_1^n$ if and only if the subtree of the left node that is labelled by a is n -bisimilar to the subtree of the right node that is labelled by a .*

Proof. The proof follows the semantic encoding of a pebble game. If the left and right a -marked subtree are n -bisimilar, then Duplicator has a winning strategy in the game. Thus if $M_s \models \mathcal{B}_1^n$, any move Spoiler makes corresponds to a refinement that makes $trunk_1$ true at one branch, and $trunk_2$ true at the other. But for any move that Spoiler makes, Duplicator may find a move (a refinement that makes $trunk_2$ true for both subtrees) where the remaining subtrees are $(n-1)$ -bisimilar (and thus \mathcal{B}_2^n is true for the refined binary tree). Applying the argument inductively, it follows that if the left and right a -marked subtrees are n -bisimilar, then $M_s \models \mathcal{B}_1^n$.

Conversely, if $M_s \models \mathcal{B}_1^n$, then we may extract a winning strategy for Duplicator in the pebble game. Any move that Spoiler may make in the game will correspond to a refinement, M_s^1 that makes $trunk_1$ true at one subtree and $trunk_2$ true at the other. As $M_s \models \mathcal{B}_1^n$, for every such refinement, there is a further refinement, M_s^2 that has both




Fig. 4. The state of the game after four moves. The winning move for Spoiler to play is to select the bold successor. Then no matter what successor Duplicator picks, Spoiler may select one of it's successors, leaving Duplicator with no move to make

a -marked subtrees satisfying $trunk_2$, and furthermore, $M_s^2 \models \mathcal{B}_2^n$. Therefore, Duplicator may chose to move according to this bisimulation (i.e. by selecting which ever successor was preserved in the refinement). As $M_s \models \mathcal{B}_1^n$, Duplicator's strategy is guaranteed to last at least n moves and thus the left and right a -marked subtrees are n -bisimilar. \square

Lemma 6.7 *No formula of the modal μ -calculus can express the property of $(n + 1)$ -bisimilarity in size less than $2^n/4$.*

Proof. We note that there are roughly 2^{2^n} non-bisimilar (1,2)-trees of height $n + 1$. The actual number is specified by the recurrence $f(n) = (f(n - 1)^2 + f(n - 1))/2$ where $f(0) = 2$ and a simple induction will show that $f(n) > 2^{2^n/4}$. As every formula of the μ -calculus is expressively equivalent to an alternating automaton of equal size, if a formula of size less than $2^n/4$ were able to express $n + 1$ -bisimilarity then an alternating automaton of size less than $2^n/4$ would be able to accept all pairs of subtrees that are $n + 1$ -bisimilar. Given the 2-player parity game that results from applying the alternating automaton to any we may associate with every distinct sub-tree (up to bisimulation) the set of automaton states for which the automaton player has a winning strategy from that state in the the subtree. As there are more than $2^{2^n/4}$ non-bisimilar subtrees and less than $2^n/4$ states, there must be two non-bisimilar subtrees, T_1 and T_2 for which an automaton winning strategy exists for exactly the same set of states. The alternating automaton would not be able to distinguish the case where T_1 is the left successor and T_2 is the right successor (which it cannot accept), and the case where T_2 is the left and right successor (which it does accept). Therefore, any formula that expresses $n + 1$ -bisimilarity for (1, 2)-trees must have at least $2^n/4$ subformulas. \square

Corollary 6.8 $\mathcal{L}_\triangleright$ is exponentially more succinct than \mathcal{L} .

Proof. From the lemmas above, $\mathcal{L}_\triangleright$ is able to express n -bisimilarity with a formula of

size $O(n^2)$, while \mathcal{L} requires a formula of size $2^n/4$ at least. The quadratic growth of the $\mathcal{L}_\triangleright$ formula is cancelled out by a constant in the exponent of the growth rate of the \mathcal{L} formula, so we may find a family of formulas in $\mathcal{L}_\triangleright$, such that for a formula of size n , the smallest equivalent \mathcal{L} formula is bound below by $2^{O(n)}$. \square

We note this proof applies without change to show $\mathcal{L}_\triangleright^\mu$ is exponentially more succinct than \mathcal{L}^μ .

7 Discussion and perspectives

The logic $\mathcal{L}_\triangleright$ is presented with respect to the class of all epistemic models. By restricting the class of models the logic is interpreted over we may associate different meanings with the modalities. For example, the epistemic logic **S5** is interpreted over all models where the accessibility relation is reflexive, transitive and symmetric (we will denote this class $S5$), and the logic **K4** is interpreted over all models with a transitive accessibility relation (denoted $K4$). Given any class of models \mathcal{C} , we define the logic $\mathcal{L}_\triangleright^{\mathcal{C}}$ to be as in Section 3 except:

- (i) The interpretation is restricted to models in the class \mathcal{C}
- (ii) The semantic interpretation of \blacktriangleright is given by:

$$M_s \models \blacktriangleright_a \phi \text{ iff for all } M_{s'} \in \mathcal{C} : M_s \succeq_a M_{s'} \text{ implies } M_{s'} \models \phi.$$

A study of how various classes of models affect the properties of bisimulation quantified logics is given in [11]. For the effect of varying classes of models on the axiomatization given, we note that while the schema **FEL** is sound for $\mathcal{L}_\triangleright$, it is not the case that the axiom **GK** is sound for restricted classes of models. For example in the class of reflexive, transitive and symmetric models (i.e. $S5$ frames) we have $\diamond\triangleright\Box p \wedge \diamond\triangleright\neg\Box p$ is consistent, but $\triangleright\nabla(\Box p, \neg\Box p)$ is not. In future work we will examine axiomatizations and complexity for refinement quantifiers in logics such as **S5**, **KD45** and **K4**.

References

- [1] Aczel, P., “Non-Well-Founded Sets,” CSLI Publications, Stanford, CA, 1988, cSLI Lecture Notes 14.
- [2] Alur, R., T. A. Henzinger and O. Kupferman, *Alternating-time temporal logic*, Journal of the ACM **49** (2002), pp. 672–713.
- [3] Arnold, A. and D. Niwinski, “Rudiments of μ -calculus,” North Holland, 2001.
- [4] Baltag, A., L. Moss and S. Solecki, *The logic of public announcements, common knowledge, and private suspicions*, in: I. Gilboa, editor, *Proceedings of the 7th Conference on Theoretical Aspects of Rationality and Knowledge (TARK 98)*, 1998, pp. 43–56.
- [5] Bilkova, M., A. Palmigiano and Y. Venema, *Proof systems for the coalgebraic cover modality.*, in: C. Areces and R. Goldblatt, editors, *Advances in Modal Logic* (2008), pp. 1–21.

- [6] Blackburn, P., M. de Rijke and Y. Venema, “Modal Logic,” Cambridge University Press, Cambridge, 2001, Cambridge Tracts in Theoretical Computer Science 53.
- [7] D’Agostino, G. and M. Hollenberg, *Logical questions concerning the mu-calculus: Interpolation, Lyndon and Los-Tarski*, Journal of Symbolic Logic **65** (2000), pp. 310–332.
- [8] D’Agostino, G. and G. Lenzi, *An axiomatization of bisimulation quantifiers via the μ -calculus*, Theor. Comput. Sci. **338** (2005), pp. 64–95.
- [9] D’Agostino, G. and G. Lenzi, *A note on bisimulation quantifiers and fixed points over transitive frames.*, J. Log. Comput. **18** (2008), pp. 601–614.
- [10] Fagin, R., J. Halpern, Y. Moses and M. Vardi, “Reasoning about Knowledge,” MIT Press, Cambridge MA, 1995.
- [11] French, T., “Bisimulation quantifiers for modal logic,” Ph.D. thesis, University of Western Australia (2006).
- [12] Harel, D., D. Kozen and J. Tiuryn, “Dynamic Logic,” MIT Press, 2000.
- [13] Janin, D. and I. Walukiewicz, *Results on the expressive completeness of the propositional mu-calculus with respect to monadic second order logic*, in: *Concurrency Theory, 7th International Conference*, LNCS **1119** (1996), pp. 263–277.
- [14] Kupferman, O., M. Vardi and P. Wolper, *Module checking*, Information and Computation **164** (2001), pp. 322–344.
- [15] Lomuscio, A. and M. Ryan, *An algorithmic approach to knowledge evolution*, Artificial Intelligence for Engineering Design, Analysis and Manufacturing (AIEDAM) **13(2)** (1998).
- [16] Morgan, C., “Programming from Specifications: Second Edition,” Prentice Hall International, Hemstead, UK, 1994.
URL <http://web.comlab.ox.ac.uk/oucl/publications/books/PFS/>
- [17] Pauly, M., “Logic for social software,” Ph.D. thesis, University of Amsterdam (2001), iLLC Dissertation Series DS-2001-10.
- [18] Ramadge, P. and W. Wonham, *On the supervisory control of discrete event systems*, in: *Proceedings of the IEEE*, 1989, pp. 81–98.
- [19] Stirling, C., *The joys of bisimulation*, in: *MFCS*, 1998, pp. 142–151.
- [20] Tsitsoklis, J., *On the control of discrete event dynamical systems*, Mathematics of Control Signals and Systems **2** (1989), pp. 95–107.
- [21] van Ditmarsch, H. and T. French, *Simulation and information*, in: J. Broersen and J.-J. Meyer, editors, *Proceedings of KRAMAS, Sydney*, LNAI **5605** (2009), pp. 51–65.
- [22] van Ditmarsch, H., W. van der Hoek and B. Kooi, “Dynamic Epistemic Logic,” Synthese Library **337**, Springer, 2007.
- [23] Walukiewicz, I., *Completeness of Kozen’s axiomatisation of the propositional mu-calculus*, INFCTRL: Information and Computation (formerly Information and Control) **157** (2000).