J. van Benthem
N. Bezhanishvili
S. Enqvist

A Propositional Dynamic Logic for Instantial Neighborhood Semantics

Abstract.   We propose a new perspective on logics of computation by combining instantial neighborhood logic INL with bisimulation safe operations adapted from PDL. INL is a recent modal logic, based on an extended neighborhood semantics which permits quantification over individual neighborhoods plus their contents. This system has a natural interpretation as a logic of computation in open systems. Motivated by this interpretation, we show that a number of familiar program constructors can be adapted to instantial neighborhood semantics to preserve invariance for instantial neighborhood bisimulations, the appropriate bisimulation concept for INL. We also prove that our extended logic IPDL is a conservative extension of dual-free game logic, and its semantics generalizes the monotone neighborhood semantics of game logic. Finally, we provide a sound and complete system of axioms for IPDL, and establish its finite model property and decidability.

Keywords: Dynamic logic, game logic, neighborhood models, instantial neighborhood logic

Author affiliations

Johan van Benthem: Institute for Logic, Language and Computation, University of Amsterdam, Netherlands.
Nick Bezhanishvili: Institute for Logic, Language and Computation, University of Amsterdam, Netherlands.
Sebastian Enqvist: Department of Philosophy, Stockholm University, Sweden.

1.   Introduction

In this paper, we introduce a new modal logic of computation, in the style of propositional dynamic logic, based on *instantial neighborhood logic* INL [6]. The logic INL is based on a recent variant of monotone neighborhood semantics for modal logics, called instantial neighborhood semantics. In the standard neighborhood semantics, the box operator has the interpretation: $\Box p$ is true at a point if *there exists* a neighborhood in which *all* the elements

satisfy the proposition p. So the box operator has a built-in fixed existential-universal quantifier pattern. In instantial neighborhood logic, we allow both universal and existential quantification over individual neighborhoods, so the basic modality has the form $\Box(p_1, ..., p_n; q)$. This formula is true at a point if *there exists* a neighborhood N in which *all* the elements satisfy the proposition q, and furthermore each of the propositions $p_1, ..., p_n$ are satisfied by *some* elements of N. INL is more expressive than monotone neighborhood logic, and comes with a natural associated notion of bisimulation together with a Hennessy-Milner theorem for finite models. It has a complete system of axioms, has the finite model property, is decidable and PSpace-complete.

Formally, our proposal is to extend the base language INL with bisimulation safe "program constructors", as in standard propositional dynamic logic of sequential programs (PDL). The usual repertoire here consists of choice, test, sequential composition and a Kleene star for program iteration. Similar additions have been studied extensively for the standard (monotone) neighborhood semantics, with constructors interpreted as ways of constructing complex *games* (this idea dates back to [19]). In the neighborhood setting, additional operations are available, including the *dual* construction. This is a very powerful device, and it is well known that dynamic game logic is not contained in any fixed level of the μ-calculus alternation hierarchy [8].

We think of our extended system of 'instantial PDL' (IPDL for short), as a dynamic logic for a richer notion of computation than sequential programs, which is sometimes referred to as *open systems* [2]. In open systems, a computational process is viewed as an agent acting in an uncertain environment that affects the outcome of each action. That is, each action by the agent is followed by a response from the environment, which is not uniquely determined. This is in contrast with *reactive systems*, where the behaviour of the system is non-deterministic but completely determined by the actions of the agent [1]. Many different logics for open systems have been proposed, perhaps the most well known being the alternating-time temporal logic ATL introduced by Alur et al. in [2]. Dynamic game logic can be interpreted in a similar way, thinking of processes as "games against the environment". Game logic is usually interpreted with a neighborhood semantics, in which neighborhoods of "worlds" in a model are taken to represent powers of some player, i.e. goals that can be enforced by some action or strategy. Instantial neighborhood semantics introduces a more fine-grained perspective to this setting, with a more expressive language and a finer bisimulation concept than standard neighborhood bisimilarity, namely the instantial neighborhood bisimulations of [6]. Since INL formulas allow existential quantification over individual neighborhoods, this language is suitable to describe not only

what conditions an agent can enforce by some action, but allows more precise reasoning about exactly what possible outcomes may result from some action. Concretely, we introduce formulas of the following kind:

$$\langle a \rangle (\psi_1, ..., \psi_n; \varphi)$$

expressing the following property about the system/program a: "the agent can act so as to ensure that φ holds, while allowing (for each $i \in \{1, ..., n\}$) the possibility that the property ψ_i may hold". In other words, instantial neighborhood logic has a natural interpretation as a simple yet expressive modal logic for computation in open systems.

However, on a computational interpretation, it is standard wisdom that one needs to extend the language to allow certain fixpoint constructions, since most specifications of systems that turn up in practice – safety, liveness, fairness etc. – involve fixpoints. There are many options available here, the most obvious one being to simply add unrestricted fixpoint operators as in the full modal μ-calculus. This route is already well understood: it was noted in [6] that INL is a *coalgebraic modal logic* in a completely standard sense, and so the μ-calculus extension of INL is a coalgebraic modal μ-calculus as in [22, 14]. Such coalgebraic μ-calculi have been quite extensively studied, with generic results on decidability and complexity, [11] and completeness [12, 13]. But there are also other versions of modal fixpoint logics, often corresponding to fragments of μ-calculi. Most notably these include propositional dynamic logics like PDL or game logic, and temporal logics like CTL or ATL. Thus an obvious point on the agenda, for further exploration of INL as a modal logic of computation, is to develop dynamic and temporal logic extensions of INL. This paper deals with the former, and sets up a propositional dynamic logic interpreted over instantial neighborhood semantics.

Overview of the paper

We first introduce syntax and semantics of instantial neighborhood logic, and its extensions leading up to the full language IPDL, provide sound and complete systems of axioms, and establish bisimulation invariance and decidability.

The completeness proof for IPDL, including all program constructors considered, is based on the standard completeness proof for PDL (see [9] for an exposition), but involves non-trivial new features. In particular, the system requires two distinct induction rules, corresponding to a nested least fixpoint induction, and the model construction makes heavy use of a normal

form for INL-formulas established in [6]. Finally, we prove that our logic is a conservative extension of the dual-free fragment of dynamic game logic.

This paper is an extended version of a conference paper presented at LORI VI 2017 [5]. The technical results presented here are the same as in that paper, but we have added full proofs. We have also added an example illustrating how the logic can be used to reason about open systems computation, and a discussion on the informal interpretation of the program operations in the language, relating this to the issue of bisimulation safety.

Acknowledgement   We thank the referee for valuable feedback and for spotting an error in the original version of the proof of Theorem 7. We also thank a number of colleagues for helpful discussions on earlier versions of the manuscript, where in particular we wish to mention Valentin Goranko, Helle Hansen, Tadeusz Litak and Lutz Schröder.

2.   Instantial neighborhood logic

2.1.   Syntax and semantics

We start by reviewing the basic language for instantial neighborhood semantics. The only difference with [6] is that we interpret the language over *labelled* neighborhood structures, where labels play the same role as "atomic programs" in PDL. The syntax of INL is given by the following grammar:

$$\varphi := p \in \mathsf{Prop} \mid \varphi \wedge \varphi \mid \neg\varphi \mid \langle a \rangle(\Psi; \varphi)$$

where a ranges over a fixed set $\mathcal{A}$ of *atomic labels*, and Ψ ranges over finite sets of formulas of INL. We deviate a bit from the syntax of [6] here in allowing Ψ to be a finite *set* rather than a tuple of formulas. We shall sometimes write $\langle a \rangle(\psi_1, ..., \psi_n; \varphi)$ rather than $\langle a \rangle(\{\psi_1, ..., \psi_n\}; \varphi)$, in particular, we write $\langle a \rangle(\psi; \varphi)$ rather than $\langle a \rangle(\{\psi\}; \varphi)$, and $\langle a \rangle\varphi$ rather than $\langle a \rangle(\emptyset; \varphi)$.

The modalities of INL have a number of interpretations. In the present setting, we interpret INL in terms of computation in open systems, so that the formula $\langle a \rangle(\psi_1, ..., \psi_n; \varphi)$ is informally interpreted as saying: "in the system a, the agent has an action to enforce the condition φ while simultaneously allowing possible outcomes satisfying each of the conditions ψ_i".

EXAMPLE 1. Consider the following example: three separate servers are shared by a number of agents and protected by passwords available to the users. Each server can only be accessed by one user at a time. Taking the perspective of one of the agents, let A_i stand for "the agent has access to

server S_i", for $i \in \{1, 2, 3\}$, and let O_i stand for "server S_i is occupied". If we introduce a name σ for the system so described, then the following is true for each given user, in each given state of the system σ:

$$\neg \langle \sigma \rangle (\neg O_1; A_1) \wedge \neg \langle \sigma \rangle (\neg O_2; A_2) \wedge \neg \langle \sigma \rangle (\neg O_3; A_3)$$

This expresses that the user cannot log in to a server without blocking the other users from having access to that server. The following also holds:

$$\neg O_3 \rightarrow \langle \sigma \rangle (\neg O_1, \neg O_2; A_3)$$

If server S_3 is available then the agent can access it while leaving servers S_1 and S_2 available to be occupied by other users. Note the distinction here: the user cannot *guarantee* that the servers S_1,S_2 will be available, they might be occupied by other users, but she can *allow* them to remain available. Finally, the following holds:

$$\neg \langle \sigma \rangle (\neg A_1, \neg A_2; (O_1 \rightarrow A_1) \vee (O_2 \rightarrow A_2))$$

This last example is perhaps less obvious: it says that the only way a user can make sure that at least one of the servers S_1 or S_2 will not be occupied by some other user is to log in to at least one of them herself.

For the formally precise semantics, formulas in INL will be interpreted over neighborhood structures.

DEFINITION 1. A *neighborhood* frame is a structure (W, R) where W is a set and R associates with each $a \in \mathcal{A}$ a binary relation $R_a \subseteq W \times \mathcal{P}W$. A *neighborhood model* (W, R, V) is a neighborhood frame together with a valuation $V : \mathsf{Prop} \rightarrow \mathcal{P}W$.

DEFINITION 2. We define the interpretations of all formulas in a neighborhood model $\mathfrak{M} = (W, R, V)$ as follows:

- $[\![p]\!] = V(p)$.

- $[\![\varphi \wedge \psi]\!] = [\![\varphi]\!] \cap [\![\psi]\!]$.

- $[\![\neg \varphi]\!] = W \setminus [\![\varphi]\!]$.

- $u \in [\![\langle a \rangle (\psi_1, ..., \psi_k; \varphi)]\!]$ iff there is some $Z \subseteq W$ such that:

$(u, Z) \in R_a$ and $Z \subseteq [\![\varphi]\!]$, $Z \cap [\![\psi_i]\!] \neq \emptyset$ for $i \in \{1, ..., k\}$

We write $\mathfrak{M}, v \Vdash \varphi$ for $v \in [\![\varphi]\!]$, and we write $\Vdash \varphi$ and say that φ is *valid* if, for every neighborhood model $\mathfrak{M}$ and $v \in W$, we have $\mathfrak{M}, v \Vdash \varphi$. We allow the notation $[\![-]\!]_{\mathfrak{M}}$ to make explicit reference to the model in the background.

Neighborhood models come with a natural notion of bisimulation, introduced in a more general setting in [6]. For this definition, the so called *Egli-Milner lifting* of a binary relation will play an important role:

DEFINITION 1. The *Egli-Milner lifting* of a binary relation $R \subseteq X \times Y$, denoted $\overline{R}$, is a relation from $\mathcal{P}X$ to $\mathcal{P}Y$ defined by: $Z\overline{R}Z'$ iff:

1. For all $z \in Z$ there is some $z' \in Z'$ such that zRz'.

2. For all $z' \in Z'$ there is some $z \in Z$ such that zRz'.

We write $R; S$ for the composition of relations R and S. It is well known that the Egli-Milner lifting preserves relation composition:

$$\overline{R; S} = \overline{R}; \overline{S}$$

DEFINITION 2. Let $\mathfrak{M} = (W, R, V)$ and $\mathfrak{M}' = (W', R', V')$ be any neighborhood models. The relation $B \subseteq W \times W'$ is said to be an *instantial neighborhood bisimulation* if for all uBu' and all atomic labels a we have:

Atomic For all p, $u \in V(p)$ iff $u' \in V'(p)$.

Forth For all Z such that uR_aZ, there is some Z' such that $u'R'_aZ'$ and $Z\overline{B}Z'$.

Back For all Z' such that $u'R'_aZ'$ there is some Z such that uR_aZ and $Z\overline{B}Z'$.

We say that pointed models $\mathfrak{M}, w$ and $\mathfrak{N}, v$ are *bisimilar*, written $\mathfrak{M}, w \leftrightarrow \mathfrak{N}, v$, if there is an instantial neighborhood bisimulation B between $\mathfrak{M}$ and $\mathfrak{N}$ such that wBv.

It is easy to check that all formulas of INL are invariant for instantial neighborhood bisimilarity:

PROPOSITION 1. *If $\mathfrak{M}, w \leftrightarrow \mathfrak{N}, v$ then $\mathfrak{M}, w \Vdash \varphi$ iff $\mathfrak{N}, v \Vdash \varphi$, for each formula φ of the language* INL.

2.2. Axiomatization

We now turn to the task of axiomatizing the valid formulas of INL. Our system of axioms is a gentle modification of the axiom system for instantial neighborhood logic presented in [6]. The axioms and rules consist of all propositional tautologies, plus the following schemas:

INL axioms

Weak: $\langle a \rangle (\Psi; \varphi) \to \langle a \rangle (\Psi'; \varphi)$ for $\Psi' \subseteq \Psi$

Un: $\langle a \rangle (\psi_1, ..., \psi_n; \varphi) \to \langle a \rangle (\psi_1 \wedge \varphi, ..., \psi_n \wedge \varphi; \varphi)$

Lem: $\langle a \rangle (\Psi; \varphi) \to \langle a \rangle (\Psi \cup \{\gamma\}; \varphi) \vee \langle a \rangle (\Psi; \varphi \wedge \neg\gamma)$

Bot: $\neg\langle a \rangle (\bot; \varphi)$

Rules

MP:
$$\frac{\varphi \to \psi \qquad \varphi}{\psi}$$

Mon:
$$\frac{\psi_1 \to \alpha_1 \ \ldots \ \psi_n \to \alpha_n \quad \varphi \to \beta}{\langle a \rangle (\psi_1, ..., \psi_n; \varphi) \to \langle a \rangle (\alpha_1, ..., \alpha_n; \beta)}$$

It is routine to derive the usual rule of replacement of equivalents:

RE:
$$\frac{\varphi \leftrightarrow \psi \qquad \theta}{\theta[\varphi/\psi]}$$

where $\theta[\varphi/\psi]$ is the result of substituting some occurrences of the formula ψ by φ in θ.

We denote this system of axioms by Ax1 and write $\text{Ax1} \vdash \varphi$ to say that the formula φ is provable in this axiom system. We also write $\varphi \vdash_{\text{Ax1}} \psi$ for $\text{Ax1} \vdash \varphi \to \psi$, and say that φ *provably entails* ψ.

THEOREM 1. *The system* Ax1 *is sound and complete for validity on neighborhood models.*

The proof of this result is essentially the same as in [6], and will not be repeated here. Since the proof in [6] constructs a finite model for each consistent formula, we also get:

THEOREM 2. *The logic* INL *has the finite model property and is decidable.*

EXAMPLE 2. Continuing from Example 1, we recall the formula:

$$\neg\langle\sigma\rangle(\neg O_i; A_i)$$

expressing that a user cannot both log in to a server and leave it available to other users. This reduces, of course, to the fact that the formula $A_i \to O_i$ is true in every state: a server cannot be both accessed by a user and at the same time not occupied. So we can take this formula instead as an extra assumption. By replacing equivalent formulas we then get the implication:

$$\langle\sigma\rangle(\neg O_i; A_i) \to \langle\sigma\rangle(\neg O_i; A_i \wedge O_i)$$

We can now apply the axiom (Un) to get the implication:

$$\langle\sigma\rangle(\neg O_i; A_i \wedge O_i) \to \langle\sigma\rangle(\neg O_i \wedge A_i \wedge O_i; A_i \wedge O_i)$$

Replacing equivalents again we get:

$$\langle\sigma\rangle(\neg O_i; A_i \wedge O_i) \to \langle\sigma\rangle(\bot; A_i \wedge O_i)$$

But as an instance of (Bot) we have the implication:

$$\langle\sigma\rangle(\bot; A_i \wedge O_i) \to \bot$$

So we get:

$$\langle\sigma\rangle(\neg O_i; A_i) \to \bot$$

i.e. $\neg\langle\sigma\rangle(\neg O_i; A_i)$ as required.

3.   Basic program operations

3.1.   Semantics and basic model theory

In what follows we shall extend the language INL with program operations, corresponding to known operations from PDL. We also include the "dual choice" constructor from dynamic game logic. Of course, there are design choices to make here, and we need to set up criteria for what counts as a correct definition of each program operation. We shall follow these three:

1. The constructions should be as simple as possible.
2. Each operation should be a natural adaptation of the corresponding operation from PDL to the INL framework, with minimal modifications.

3. Most importantly: each operation should be *bisimulation safe*, i.e. the dynamic logic extending INL with all the program operations should remain invariant for instantial neighborhood bisimulations.

We first extend the language INL with four basic PDL-style operations: test, choice, parallel composition and sequential composition. The resulting language will be called *dynamic instantial neighborhood logic*, or (DINL). The syntax of DINL is defined by the following dual grammar.

$$\varphi := p \in \mathsf{Prop} \mid \varphi \wedge \varphi \mid \neg\varphi \mid \langle \pi \rangle (\Psi; \varphi)$$

$$\pi := a \in \mathcal{A} \mid \varphi? \mid \pi \cup \pi \mid \pi \cap \pi \mid \pi;\pi$$

The operation $\cup$ is interpreted as non-deterministic choice between two programs for the agent: $\pi_1 \cup \pi_2$ means "either do π_1 or do π_2". The operation $\cap$ is intepreted as a choice between two programs for the environment: $\pi_1 \cap \pi_2$ means "do π_1 and π_2 in parallell". Formally, the operation $\cap$ is similar to the parallell composition in concurrent PDL (see [15]). Finally, the operator ; is interpreted as sequential composition: $\pi_1;\pi_2$ means "first do π_1 then do π_2". We define the formal interpretation $[\![o]\!]$ of each operation $o \in \{\cup, \cap, ;\}$ in a neighborhood model $\mathfrak{M}$ as a binary map from pairs of neighborhood relations to neighborhood relations, as follows:

– $R_1 [\![\cup]\!] R_2 = R_1 \cup R_2$
– $R_1 [\![\cap]\!] R_2 = \{(w, Z_1 \cup Z_2) \mid (w, Z_1) \in R_1 \ \& \ (w, Z_2) \in R_2\}$
– $(w, Z) \in R_1 [\![;]\!] R_2$ iff there is some set Y and some family of sets F such that $(w, Y) \in R_1$, $(Y, F) \in \overline{R}_2$ and $Z = \bigcup F$.

The interpretation $[\![?]\!]$ of the test operator will be a map $[\![?]\!]$ assigning a neighborhood relation to each subset Z of W, defined by:

$$[\![?]\!]Z := \{(u, \{u\}) \mid u \in Z\}$$

We defer a more detailed discussion of the informal interpretation of the program operations to Section 3.2. Note that $[\![?]\!]$ is monotone in the sense that $Z \subseteq Z'$ implies $[\![?]\!]Z \subseteq [\![?]\!]Z'$. Each operator $o \in \{\cup, \cap, ;\}$ is also monotone, in the sense that $R_1 [\![o]\!] R_2 \subseteq R_1' [\![o]\!] R_2'$ whenever $R_1 \subseteq R_1'$ and $R_2 \subseteq R_2'$. For the sequential composition operator, this uses the well known fact that the Egli-Milner lifting is monotone, i.e. $\overline{R} \subseteq \overline{R'}$ whenever $R \subseteq R'$.

DEFINITION 3. Given a neighborhood model, we define the semantic interpretations of all formulas, and the neighborhood relations corresponding to

all complex labels π, by the following mutual recursion:

- $\llbracket p \rrbracket = V(p)$.

- $\llbracket \varphi \wedge \psi \rrbracket = \llbracket \varphi \rrbracket \cap \llbracket \psi \rrbracket$.

- $\llbracket \neg \varphi \rrbracket = W \setminus \llbracket \varphi \rrbracket$.

- $u \in \llbracket \langle \pi \rangle (\psi_1, ..., \psi_k; \varphi) \rrbracket$ iff there is some $Z \subseteq W$ such that:

$$(u, Z) \in R_\pi \text{ and } Z \subseteq \llbracket \varphi \rrbracket,\ Z \cap \llbracket \psi_i \rrbracket \neq \emptyset \text{ for } i \in \{1, ..., k\}.$$

- $R_{\pi_1 o \pi_2} = R_{\pi_1} \llbracket o \rrbracket R_{\pi_2}$ for $o \in \{\cup, \cap, ;\}$.

- $R_{\varphi?} = \llbracket ? \rrbracket \llbracket \varphi \rrbracket$

The definitions of the dynamic operations stated abobve are tailored towards obtaining the following result:

PROPOSITION 2. *All formulas of* DINL *are invariant for instantial neighborhood bisimulations.*

PROOF. We first prove the following claim, expressing bisimulation safety of the operations that we have introduced:

CLAIM 1. Let B be an instantial neighborhood bisimulation between models $\mathfrak{M} = (W, R, V)$ and $\mathfrak{M}' = (W', R', V')$. Then for any complex label π, such that every term of the form $\varphi?$ appearing in π, φ is invariant for instantial neighborhood bisimulations, and for any $u \in W$ and $u' \in W'$ with uBu':

Forth For all Z such that $uR_\pi Z$, there is some Z' such that $u'R'_\pi Z'$ and $Z\overline{B}Z'$.

Back For all Z' such that $u'R'_\pi Z'$ there is some Z such that $uR_\pi Z$ and $Z\overline{B}Z'$.

We prove the Claim by induction on the complexity of labels. For atomic labels the result holds by definition. For the inductive steps, we only prove the "Forth" clause, as the "Back" clause follows by a symmetric argument. For the test operator, the result follows immediately from the assumption that every formula appearing in a sub-term of π is bisimulation invariant.

For choice, suppose $(u, Z) \in R_{\pi_1 \cup \pi_2}$. Then $(u, Z) \in R_{\pi_1}$ or $(u, Z) \in R_{\pi_2}$, say, the first holds. Then by the Forth clause for π_1 there is some Z' with $(u', Z') \in R'_{\pi_1}$ such that $Z\overline{B}Z'$. Since $(u', Z') \in R'_{\pi_1 \cup \pi_2}$ also, we are done.

For dual choice, suppose $(u, Z) \in R_{\pi_1 \cap \pi_2}$. Then $Z = Z_1 \cup Z_2$ where $(u, Z_1) \in R_{\pi_1}$ and $(u, Z_2) \in R_{\pi_2}$. By the Forth condition for π_1 and π_2 we find sets Z_1' and Z_2' such that $(u', Z_1') \in R_{\pi_1}'$, $(u', Z_2') \in R_{\pi_2}'$ and $Z_1 \overline{B} Z_1'$, $Z_2 \overline{B} Z_2'$. We leave it to the reader to check that:

$$(Z_1 \cup Z_2, Z_1' \cup Z_2') \in \overline{B}.$$

Since $(u', Z_1' \cup Z_2') \in R_{\pi_1 \cap \pi_2}'$, we are done.

Finally, for sequential composition, suppose there is a set X such that $(u, X) \in R_{\pi_1 ; \pi_2}$, witnessed by a set Y such that $(u, Y) \in R_{\pi_1}$ and a family $F \subseteq \mathcal{P}(W)$ such that $(Y, F) \in \overline{R}_{\pi_2}$ and $X = \bigcup F$. By the Forth condition for π_1 there is a set Y' such that $(u', Y') \in R_{\pi_1}'$ and $Y \overline{B} Y'$. We define a family $F' \subseteq \mathcal{P}(W')$ as follows: set $Z' \in F'$ iff there is some $v' \in Y'$, some $v \in Y$ and some $Z \in F$ such that: $(v', Z') \in R_{\pi_2}'$, $(v, Z) \in R_{\pi_2}$ and $Z \overline{B} Z'$.

First, we claim that $(Y', F') \in \overline{R'}_{\pi_2}$: first, if $Z' \in F'$ then it is immediate from the definition that $(v', Z') \in R_{\pi_2}'$ for some $v' \in Y'$. Conversely, given $v' \in Y'$, since $Y \overline{B} Y'$ there must be some $v \in Y$ with vBv', and since $(Y, F) \in \overline{R}_{\pi_2}$ there is some $Z \in F$ with $(v, Z) \in R_{\pi_2}$. But then, by the Forth condition for π_2 there must be some Z' with $(v', Z') \in R_{\pi_2}'$ and $Z \overline{B} Z'$. We immediately get $Z' \in F'$, as required.

We now show that:
$$\left(\bigcup F, \bigcup F' \right) \in \overline{B}.$$

To see this, suppose first that $w \in \bigcup F$. Then $w \in Z$ for some $Z \in F$. Since $(Y, F) \in \overline{R}_{\pi_2}$ there is some $v \in Y$ with $(v, Z) \in R_{\pi_2}$. Since $Y \overline{B} Y'$ there is some $v' \in Y'$ such that vBv'. By the Forth condition for π_2 there is some Z' with $Z \overline{B} Z'$ and $(v', Z') \in R_{\pi_2}'$. We get $Z' \in F'$, and there must be some $w' \in Z'$ with wBw'. But then $w' \in \bigcup F'$, as required.

Conversely, suppose $w' \in \bigcup F'$. Then $w' \in Z'$ for some $Z' \in F'$. By definition of F', there is a $Z \in F$ with $Z \overline{B} Z'$, and so there is some $w \in Z$ with wBw'. But then $w \in \bigcup F$ as required, and the claim is proved.

The proposition now follows from the claim by a routine argument. ∎

3.2. Informal interpretation

The neighborhood relation R_π associated with a program term π in a neighborhood model $\mathfrak{M}$ should be understood as follows: at each point w in a model, there is a certain family of available *actions of type π* that the agent can perform. Each such action α corresponds to a neighborhood $Z \in R_\pi[w]$, and Z represents the possible outcomes of the action α, as determined by the response of the environment. The interpretations of choice $\cup$ and dual

choice $\cap$ should thus be clear: an action of type $\pi_1 \cup \pi_2$ is simply an action of either type π_1 or π_2, and so the definition of $[\![\cup]\!]$ as union of neighborhood relations is the natural one. For dual choice, an action α of type $\pi_1 \cap \pi_2$ consists of an action β_1 of type π_1 and an action β_2 of type π_2, where the action actually performed is determined by the environment. So a possible outcome of the action α is either a possible outcome of β_1 or one of β_2. This directly leads to the formal interpretation $[\![\cap]\!]$ of $\cap$ as it has been defined. The interpretation of the test operator is a straightforward adaption of the usual PDL-definition, and motivated in the same manner.

The less straightforward case is the sequential composition operation. Initially it seems clear what an action of type $\pi_1;\pi_2$ at a given state w should be: it is an action β_1 of type π_1 followed by an action β_2^v of type π_2 performed at each possible outcome state v of the action β_1 at w. A possible outcome of such an action α at w should then be an outcome of one of the actions β_2^v, where v is a possible outcome of the first action β_1. With this interpretation, one would expect the following definition, setting $(w, Z) \in R_1 [\![;]\!] R_2$ iff there is some set Y and a *function* $S : Y \to \mathcal{P}W$ such that:

1. $(w, Y) \in R_1$,
2. $(v, S_v) \in R_2$ for each $v \in Y$, and
3. $Z = \bigcup_{v \in Y} S_v$.

The conditions used in our actual definition of $[\![;]\!]$ are weaker than this, essentially allowing the assigment S to be a *relation* rather than a function. The reason we cannot use the stricter version of the composition operation is due to a technical fact: the "functional" version of the sequential composition operation *violates bisimulation safety*! The example shown in Figure 1, displaying two bisimilar rooted models, explains this.


Figure 1. Failure of bisimulation safety

In Figure 1, points are represented by bullets, neighborhoods are represented by ellipses, the dashed lines represent the neighborhood relation R_1 and the dotted lines represent R_2. In the model to the right, the root has a neighborhood $\{a, b\}$ according to the functional composition of R_1 and R_2, but not in the left model. Note that according to our "relational" definition of sequential composition, $\{a, b\}$ is a neighborhood in both models.

A possible response to this problem would be to modify our notion of instantial neighborhood bisimulation so as to recover safety. However, this route does not seem attractive, as instantial neighborhood bisimulations provide the natural bisimulation concept for INL, the basis for our dynamic logic. One could let the technical point settle the matter: bisimulation safety seems to be a minimal requirement for compositional reasoning about behaviour of systems, and our sequential composition operator recovers bisimulation safety arguably in a simple and mathematically natural way. But we believe there is no need for such a purely technical motivation: properly understood the sequential composition operator we have proposed fits well with its intended interpretation, and with the idea of open systems in general.

The idea is that the behaviour of an agent interacting with a system may depend not only on the state of the system itself, but also on other parameters: the internal state of the agent itself for example, or the state of other processes that the agent is also interacting with. This feature of taking into account possible interactions with unspecified, "external" systems is part of the motivation behind existing logics for open systems like ATL (see [2]). For example, looking back to Example 1, the state of the system itself merely specifies which of the three servers are occupied by which agent. The internal state of each agent – which in this case is a human – may for example involve the agent's current state of knowledge, preferences, intentions etc. So when we consider the actions available to one of the agents, we think of both the system and each of the agents as starting in a given "initial" state which may change through the course of the computation, for example due to communication between the agents, or interaction with other systems.

Therefore, in the special case of a composite action of the type $\pi_1;\pi_2$ executed at some state w, the action of the agent in the computation π_2 at a later state v resulting as the outcome of the computation π_1 might not be determined uniquely by the state v of the system. It may also depend on other parameters, which might change during the execution of π_1. This accounts for the extra non-determinism involved in our sequential composition operator, where an action of type $\pi_1;\pi_2$ need not specify a unique response to each outcome of the first action.

3.3.   Axiomatization

Our axiom system for DINL takes the sound and complete axioms for INL as its base, and extends it with reduction axioms for the test, choice, parallel composition and sequential composition operators. The axioms and rules are listed below; note that the INL axioms and the axioms for frame constraints are now stated for arbitrary complex labels π rather than just atoms a.

INL axioms:

(Weak), (Un), (Lem) and (Bot)

Reduction axioms:

Test: $\langle \gamma? \rangle (\Psi; \varphi) \leftrightarrow \gamma \wedge \bigwedge \Psi \wedge \varphi$

Ch: $\langle \pi_1 \cup \pi_2 \rangle (\Psi; \varphi) \leftrightarrow \langle \pi_1 \rangle (\Psi; \varphi) \vee \langle \pi_2 \rangle (\Psi; \varphi)$

Pa: $\langle \pi_1 \cap \pi_2 \rangle (\Psi; \varphi) \leftrightarrow \bigvee \{ \langle \pi_1 \rangle (\Theta_1; \varphi) \wedge \langle \pi_2 \rangle (\Theta_2; \varphi) \mid \Psi = \Theta_1 \cup \Theta_2 \}$

Cmp: $\langle \pi_1; \pi_2 \rangle (\psi_1, ..., \psi_n; \varphi) \leftrightarrow \langle \pi_1 \rangle (\langle \pi_2 \rangle (\psi_1; \varphi), ..., \langle \pi_2 \rangle (\psi_n; \varphi); \langle \pi_2 \rangle \varphi)$

Rules:

(MP) and (Mon)

We denote this system by Ax2 and write $\text{Ax2} \vdash \varphi$ to say that formula φ is provable in this axiom system. We also write $\varphi \vdash_{\text{Ax2}} \psi$ for $\text{Ax2} \vdash \varphi \rightarrow \psi$. We shall sometimes drop the reference to Ax2 to keep notation cleaner.

PROPOSITION 3 (Soundness). *If* $\text{Ax2} \vdash \varphi$*, then the formula* φ *is true on all neighborhood models.*

PROOF. We consider only the new reduction axioms. Soundness of (Ch) is immediate by the definition of $[\![\cup]\!]$, so we focus on (Test), (Pa) and (Cmp).

For (Test), let $\mathfrak{M}, u \Vdash \langle \gamma? \rangle (\Psi; \varphi)$. Then $(u, \{u\}) \in R_{\gamma?}$, which means that $\mathfrak{M}, u \Vdash \gamma$, and $\{u\} \subseteq [\![\varphi]\!]$ and $\{u\} \cap [\![\psi]\!] \neq \emptyset$ for each $\psi \in \Psi$, and hence $\mathfrak{M}, u \Vdash \bigwedge \Psi \wedge \varphi$. So $\mathfrak{M}, u \Vdash \gamma \wedge \bigwedge \Psi \wedge \varphi$. The converse is similar.

For (Pa), suppose that $\mathfrak{M}, w \Vdash \langle \pi_1 \cap \pi_2 \rangle (\Psi; \varphi)$. Then there is some set Z such that $(w, Z) \in R_{\pi_1 \cap \pi_2}$, $Z \subseteq [\![\varphi]\!]$ and $Z \cap [\![\psi]\!] \neq \emptyset$ for all $\psi \in \Psi$. Hence Z is of the form $Z_1 \cup Z_2$ where $(w, Z_1) \in R_{\pi_1}$ and $(w, Z_2) \in R_{\pi_2}$. Let $\Theta_1 = \{ \psi \in \Psi \mid Z_1 \cap [\![\psi]\!] \neq \emptyset \}$, and let $\Theta_2 = \{ \psi \in \Psi \mid Z_2 \cap [\![\psi]\!] \neq \emptyset \}$. Then, since $Z = Z_1 \cup Z_1$, we have $\Psi = \Theta_1 \cup \Theta_2$. Furthermore, we get

$$\mathfrak{M}, w \Vdash \langle \pi_1 \rangle (\Theta_1; \varphi) \wedge \langle \pi_2 \rangle (\Theta_2; \varphi)$$

as required. The converse direction of (Pa) is proved in a similar manner.

Next, we consider the case of sequential composition. For one direction of the equivalence, suppose that $\mathfrak{M}, w \Vdash \langle \pi_1; \pi_2 \rangle(\psi_1, ..., \psi_n; \varphi)$. Then there is some set Z with $(w, Z) \in R_{\langle \pi_1; \pi_2 \rangle}$, $Z \subseteq [\![\varphi]\!]$ and $Z \cap [\![\psi_i]\!] \neq \emptyset$ for each ψ_i. By definition of the composition operator, we find a set Y with $(w, Y) \in R_{\pi_1}$ and a family of sets F such that $(Y, F) \in \overline{R}_{\pi_2}$ and $Z = \bigcup F$. So for each $v \in Y$ there is some $S_v \in F$ with $(v, S_v) \in R_{\pi_2}$, and we get $S_v \subseteq [\![\varphi]\!]$ so $\mathfrak{M}, v \Vdash \langle \pi_2 \rangle \varphi$. Also, for each ψ_i there is some $S_i \in F$ with $S_i \cap [\![\psi_i]\!] \neq \emptyset$, and there must be some $v \in Y$ with $(v, S_i) \in R_{\pi_2}$, hence $\mathfrak{M}, v \Vdash \langle \pi_2 \rangle(\psi_i; \varphi)$. It follows that $\mathfrak{M}, w \Vdash \langle \pi_1 \rangle(\langle \pi_2 \rangle(\psi_1; \varphi), ..., \langle \pi_2 \rangle(\psi_n; \varphi); \langle \pi_2 \rangle \varphi)$ as required.

Conversely, suppose that $\mathfrak{M}, w \Vdash \langle \pi_1 \rangle(\langle \pi_2 \rangle(\psi_1; \varphi), ..., \langle \pi_2 \rangle(\psi_n; \varphi); \langle \pi_2 \rangle \varphi)$. Then there is some set Y such that $(w, Y) \in R_{\pi_1}$, $Y \subseteq [\![\langle \pi_2 \rangle \varphi]\!]$ and $Y \cap [\![\langle \pi_2 \rangle(\psi_i; \varphi)]\!] \neq \emptyset$ for each $i \in \{1, ..., n\}$. Let:

$$F := \{Z \subseteq W \mid Z \subseteq [\![\varphi]\!] \ \& \ (v, Z) \in R_{\pi_2} \text{ for some } v \in Y\}$$

Since $Y \subseteq [\![\langle \pi_2 \rangle \varphi]\!]$ it follows that $(Y, F) \in \overline{R}_{\pi_2}$, so $(w, \bigcup F) \in R_{\pi_1; \pi_2}$. Next, since $Y \cap [\![\langle \pi_2 \rangle(\psi_i; \varphi)]\!] \neq \emptyset$ for each $i \in \{1, ..., n\}$ it follows that $\bigcup F \cap [\![\psi_i]\!] \neq \emptyset$ for each $i \in \{1, ..., n\}$. We get $\mathfrak{M}, w \Vdash \langle \pi_1; \pi_2 \rangle(\psi_1, ..., \psi_n; \varphi)$ as required. ∎

By applying soundness of the reduction axioms, we can use a standard argument to obtain for every consistent formula φ of DINL a provably (and hence semantically) equivalent formula φ^t in INL, which is then satisfiable by Theorem 1. E.g., the formula $\langle \gamma? \rangle(\psi_1, ..., \psi_n; \varphi)^t$ ibecomes $\gamma^t \wedge \psi_1^t \wedge ... \wedge_n^t \wedge \varphi$.

THEOREM 3 (Completeness). *A formula φ of* DINL *is valid on all neighborhood models iff* Ax2 $\vdash \varphi$.

Also, the finite model property and decidability carry over from INL:

THEOREM 4. *The logic* DINL *has the finite model property and is decidable.*

4. Program iteration and the language IPDL

We now introduce the final operation that we consider here, a Kleene star for finite iteration. This operation generalizes the game iteration operation from game logic. The corresponding language will be denoted by IPDL, read "instantial PDL", and is given by the following dual grammar:

$$\varphi := p \in \mathsf{Prop} \mid \varphi \wedge \varphi \mid \neg \varphi \mid \langle \pi \rangle(\Psi; \varphi)$$

$$\pi := a \in \mathcal{A} \mid \varphi? \mid \pi \cup \pi \mid \pi \cap \pi \mid \pi; \pi \mid \pi^*$$

The operation $(-)^*$ is interpreted as finite iteration: π^* means "repeat π a finite number of times". More specifically, we think consider *action of type* π^* to be a long term strategy of the agent, such that each possible execution of this strategy consists of finitely many actions of type π.

For the formal semantic interpretation of the Kleene star, it will be useful to first define the relation skip by:

$$\mathsf{skip} := \{(w, \{w\}) \mid w \in W\}$$

DEFINITION 4. We define a relation $R^{[\xi]}$ for each ordinal ξ by induction:

— $R^{[0]} = \emptyset$

— $R^{[\xi+1]} = \mathsf{skip}[\![\cup]\!](R[\![;]\!]R^{[\xi]})$

— $R^\kappa = \bigcup_{\xi < \kappa} R^{[\xi]}$ if κ is a limit ordinal.

We define $[\![*]\!]R$ as $R^{[\xi]}$, with ξ the smallest ordinal satisfying $R^{[\xi]} = R^{[\xi+1]}$.

It is easy to see that this is a standard least fixpoint construction – and in particular, we have:

PROPOSITION 4. *Let W be a finite set and $R \subseteq W \times \mathcal{P}(W)$. Then:*

$$[\![*]\!]R = \bigcup_{n \in \omega} R^{[n]}$$

Proposition 4 does not hold for arbitrary models: unlike for PDL, the closure ordinal of the least fixpoint for the Kleene star may appear above ω. yet this does not contradict the reading of the Kleene star as finite iteration. The situation is analogous to the case of the μ-calculus formula:

$$\mu x.\square x$$

which can be thought of as expressing that "all computations are finite". It is well known that the closure ordinal of the least fixpoint of this formula can be higher than ω, which just means the formula may be true although the statement "all computations have length $\leq k$" is false for all k. Similarly, the formula $\langle \pi^* \rangle \varphi$ expresses that the condition φ can be forced by an action that only ever produces finitely many computations of type π, while there may be no finite upper bound on the number of iterations of π required.

DEFINITION 5. The semantics of IPDL-formulas in a neighborhood model $\mathfrak{M} = (W, R, V)$ is given by the following inductive clauses:

- $[\![p]\!] = V(p)$.

- $[\![\varphi \wedge \psi]\!] = [\![\varphi]\!] \cap [\![\psi]\!]$.

- $[\![\neg\varphi]\!] = W \setminus [\![\varphi]\!]$.

- $u \in [\![\langle\pi\rangle(\psi_1, ..., \psi_k; \varphi)]\!]$ iff there is some $Z \subseteq W$ such that:

  $(u, Z) \in R_\pi$ and $Z \subseteq [\![\varphi]\!]$, $Z \cap [\![\psi_i]\!] \neq \emptyset$ for $i \in \{1, ..., k\}$.

- $R_{\pi_1 o \pi_2} = R_{\pi_1}[\![o]\!]R_{\pi_2}$ for $o \in \{\cup, \cap, ;\}$.

- $R_{\varphi?} = [\![?]\!][\![\varphi]\!]$.

- $R_{\pi^*} = [\![*]\!]R_\pi$.

PROPOSITION 5. *All formulas of* IPDL *are invariant for instantial neighborhood bisimulations.*

The proof of this rests on a bisimulation safety argument, and the step for the Kleene star involves using the bisimulation safety of union and sequential composition to prove the appropriate back-and-forth conditions for each approximant $R_\pi^{[\xi]}$ of the least fixpoint $R_{\pi^*} = [\![*]\!]R_\pi$. We omit the details.

4.1. Axiomatization

Our axiomatization for IPDL is given below.

INL axioms:

(Weak), (Un), (Lem) and (Bot).

Reduction axioms from DINL:

(Test), (Ch), (Pa) and (Cmp).

Basic rules:

(MP) and (Mon).

Kleene star

Finally we add axioms and rules for iteration. The Kleene star is a least fixpoint construction, and a standard approach to axiomatizing least fixpoints is to use one *fixpoint axiom* and one *induction rule* (see [17]).

The fixpoint axiom **Fix** is stated as follows:

$$\langle \pi^* \rangle (\Psi; \varphi) \leftrightarrow (\bigwedge \Psi \wedge \varphi) \vee \langle \pi; \pi^* \rangle (\Psi; \varphi)$$

We will actually need *two* induction rules:

Ind1:

$$\frac{\varphi \to \gamma \qquad \langle \pi \rangle \gamma \to \gamma}{\langle \pi^* \rangle \varphi \to \gamma}$$

Ind2:

$$\frac{(\psi \wedge \varphi) \to \gamma \qquad \langle \pi \rangle (\gamma; \langle \pi^* \rangle \varphi) \to \gamma}{\langle \pi^* \rangle (\psi; \varphi) \to \gamma}$$

REMARK 1. The reason that we require two distinct induction rules can be seen as follows: the reduction axioms for IPDL should be interpreted as encoding a recursive translation of the language IPDL into the modal μ-calculus (interpreted on instantial neighborhood models). When we pass by formulas involving the Kleene-star in this translation, the translation will not surprisingly involve least fixpoint operators, and the induction rules then correspond to the Kozen-Park induction rules for least fixpoint operators. This step of the translation is trickier than the step for the Kleene star in a translation of PDL into the μ-calculus (see [10]), and requires use of nested least fixpoint variables. To illustrate, if a is an atomic game term and p, q are propositional variables, then the formula $\langle a^* \rangle (p; q)$ translates to:

$$\mu x.(p \wedge q) \vee \langle a \rangle (x; \mu y.q \vee \langle a \rangle y)$$

Note however that the fixpoint variables here are nested in a "weak" sense: the variable y occurs inside the scope of the outer fixpoint variable x, but is independent of it in the sense that there is no free occurrence of x in the scope of the variable binder μy.

Note also that the second induction axiom only involves a single instantial formula ψ. This is because we can "pre-process" an arbitrary formula $\langle \pi^* \rangle (\psi_1, ..., \psi_n; \varphi)$ by applying the axiom (Fix), and then the composition axiom (Cmp) to the formula $\langle \pi; \pi^* \rangle (\psi_1, ..., \psi_n; \varphi)$ to obtain the formula:

$$\langle \pi \rangle (\langle \pi^* \rangle (\psi_1; \varphi), ..., \langle \pi^* \rangle (\psi_n; \varphi); \langle \pi^* \rangle \varphi)$$

where each occurrence of $\langle \pi^* \rangle$ is followed by at most one instantial formula.

We denote this axiom system as Ax3 and write $\varphi \vdash_{\mathsf{Ax3}} \psi$ for $\mathsf{Ax3} \vdash \varphi \to \psi$. We will also sometimes drop the index Ax3, simply writing $\vdash \varphi$ or $\varphi \vdash \psi$.

THEOREM 5. *The axiom system* Ax3 *is sound and complete for validity over neighborhood models.*

We begin by checking soundness:

PROPOSITION 6 (Soundness). *If φ is provable in* Ax3 *then it is valid over all neighborhood models.*

PROOF. We focus on proving soundness of the two induction rules. For the first induction rule, suppose that the formulas $\varphi \to \gamma$ and $\langle \pi \rangle \gamma \to \gamma$ are valid. Suppose that $\mathfrak{M}, u \Vdash \langle \pi^* \rangle \varphi$. Then there is some Z such that $(u, Z) \in R_{\pi^*}$ and $Z \subseteq \llbracket \varphi \rrbracket$. By definition of R_{π^*} it suffices to prove, by induction on an ordinal ξ, that for all u, Z: if $(u, Z) \in R_\pi^{[\xi]}$ and $Z \subseteq \llbracket \varphi \rrbracket$ then $u \in \llbracket \gamma \rrbracket$. For $\xi = 0$ this is trivial, since $R_\pi^0 = \emptyset$. For a successor ordinal $\xi + 1$, if $(u, Z) \in R_\pi^{[\xi+1]}$ then either $Z = \{u\}$ or there is a set Y and a family of sets F such that $(Y, F) \in \overline{R_\pi^{[\xi]}}$, $(u, Y) \in R_\pi$ and $\bigcup F \subseteq \llbracket \varphi \rrbracket$. In the first case we get $\mathfrak{M}, u \Vdash \varphi$, hence $\mathfrak{M}, u \Vdash \gamma$. In the second case it follows that there is, for each $v \in Y$, some Z_v such that $(v, Z_v) \in R_\pi^{[\xi]}$ and $Z_v \subseteq \bigcup F \subseteq \llbracket \varphi \rrbracket$. By the induction hypothesis we get $Y \subseteq \llbracket \gamma \rrbracket$. But then $\mathfrak{M}, u \Vdash \langle \pi \rangle \gamma$, hence $\mathfrak{M}, u \Vdash \gamma$ as required. Finally, the induction step for limit ordinals is almost immediate, by the definition of $R_\pi^{[\xi]}$ as the union of all $R_\pi^{[\rho]}$ for $\rho < \xi$.

For the second induction rule, suppose that the formulas $(\psi \wedge \varphi) \to \gamma$ and $\langle \pi \rangle (\gamma; \varphi) \to \gamma$ are valid. Suppose that $\mathfrak{M}, u \Vdash \langle \pi^* \rangle (\psi; \varphi)$. Then there is some Z such that $(u, Z) \in R_{\pi^*}$ and $Z \subseteq \llbracket \varphi \rrbracket$, $Z \cap \llbracket \psi \rrbracket \neq \emptyset$. By definition of R_{π^*} it suffices to prove, by induction on an ordinal ξ, that for all u, Z: if $(u, Z) \in R_\pi^{[\xi]}$ and $Z \subseteq \llbracket \varphi \rrbracket$, $Z \cap \llbracket \psi \rrbracket \neq \emptyset$ then $u \in \llbracket \gamma \rrbracket$. For $\xi = 0$ this is trivial, since $R_\pi^0 = \emptyset$. For a successor ordinal $\xi + 1$, if $(u, Z) \in R_\pi^{[\xi+1]}$ then either $Z = \{u\}$ or there is a set Y and a family of sets F such that $(Y, F) \in \overline{R_\pi^{[\xi]}}$, $(u, Y) \in R_\pi$ and $\bigcup F \subseteq \llbracket \varphi \rrbracket$. In the first case we get $\mathfrak{M}, u \Vdash \psi \wedge \varphi$, and therefore $\mathfrak{M}, u \Vdash \gamma$. In the second case it follows that there is, for each $v \in Y$, some $Z_v \in F$ such that $(v, Z_v) \in R_\pi^{[\xi]}$ and $Z_v \subseteq \bigcup F \subseteq \llbracket \varphi \rrbracket$. Furthermore, there is some set $Z' \in F$ such that $Z' \cap \llbracket \psi \rrbracket \neq \emptyset$, and $Z' \subseteq \bigcup F \subseteq \llbracket \varphi \rrbracket$. Since $(Y, F) \in \overline{R_\pi^{[\xi]}}$ there must be some $w \in Y$ with $(w, Z') \in R_\pi^{[\xi]}$, and so, by the induction hypothesis we get $w \Vdash \gamma$. But then $Y \subseteq \llbracket \langle \pi^* \rangle \varphi \rrbracket$ (since $R_\pi^{[\xi]} \subseteq R_{\pi^*}$) and $Y \cap \llbracket \gamma \rrbracket \neq \emptyset$, so $\mathfrak{M}, u \Vdash \langle \pi \rangle (\gamma; \varphi)$. It follows that $\mathfrak{M}, u \Vdash \gamma$

as required. Finally, the induction step for limit ordinals is again immediate, by the definition of $R_\pi^{[\xi]}$ as the union of all $R_\pi^{[\rho]}$ for $\rho < \xi$. ∎

For the completeness proof, we rely heavily on the following lemma, proved in a slightly different version in [6]: fix a finite and subformula closed set of formulas Σ. An *atom* over Σ is a maximal consistent subset of Σ, and we denote the set of atoms over Σ by $\mathsf{At}(\Sigma)$. Given any atom $w \in \mathsf{At}(\Sigma)$, let $\widehat{w}$ be its conjunction, and let $\widehat{Z} = \{\widehat{w} \mid w \in Z\}$ for a set of atoms Z.

LEMMA 1. *Let* $\langle \pi \rangle (\Psi; \varphi)$ *be any formula such that each formula in* $\Psi \cup \{\varphi\}$ *is a boolean combination of formulas in* Σ. *Then* $\langle \pi \rangle (\Psi; \varphi)$ *is provably equivalent to a disjunction of formulas of the form* $\langle \pi \rangle (\widehat{Z}; \bigvee \widehat{Z})$ *for* $Z \subseteq \mathsf{At}(\Sigma)$ *being some set of atoms with* $w \vdash \varphi$ *for each* $w \in Z$ *and for all* $\psi \in \Psi$ *there is some* $v \in Z$ *with* $v \vdash \psi$.

PROOF. The required argument is very similar to [6]. ∎

We shall also need an adapted concept of Fischer-Ladner closure:

DEFINITION 3. A set Σ of formulas is said to be *Fischer-Ladner closed* if the following clauses hold:

- If $\varphi \in \Sigma$, and the main connective of φ is not $\neg$, then the formula $\neg\varphi$ is in Σ.

- Any subformula of a formula in Σ is in Σ.

- If $\langle \gamma? \rangle (\Psi; \varphi)$ is in Σ then so is $\gamma \wedge \bigwedge \Psi \wedge \varphi$.

- If $\langle \pi_1; \pi_2 \rangle (\psi_1, ..., \psi_n; \varphi) \in \Sigma$, then $\langle \pi_1 \rangle (\langle \pi_2 \rangle (\psi_1; \varphi), ..., \langle \pi_1 \rangle (\psi_n; \varphi); \langle \pi_2 \rangle \varphi)$ is in Σ too.

- If $\langle \pi_1 \cup \pi_2 \rangle (\Psi; \varphi) \in \Sigma$ then $\langle \pi_1 \rangle (\Psi; \varphi) \vee \langle \pi_2 \rangle (\Psi; \varphi) \in \Sigma$ too.

- If $\langle \pi_1 \cap \pi_2 \rangle (\Psi; \varphi) \in \Sigma$ then the formula:

$$\bigvee \{\langle \pi_1 \rangle (\Theta_1; \varphi) \wedge \langle \pi_2 \rangle (\Theta_2; \varphi) \mid \Psi = \Theta_1 \cup \Theta_2\}$$

  is in Σ too.

- If $\langle \pi^* \rangle (\Psi; \varphi) \in \Sigma$ then $(\bigwedge \Psi \wedge \varphi) \vee \langle \pi; \pi^* \rangle (\Psi; \varphi)$ is in Σ too.

LEMMA 2. *Every formula* φ *is a member of some finite Fischer-Ladner closed set of formulas.*

PROOF. The proof for this result is standard, see for example [9]. ∎

LEMMA 3. *Let Z be a set of atoms in $\mathsf{At}(\Sigma)$ and let θ be any formula (not necessarily in Σ). Then we have $\theta \vdash \bigvee \widehat{Z}$ if, and only if, every atom that is consistent with θ is also consistent with $\bigvee \widehat{Z}$.*

PROOF. The direction from left to right is trivial. From right to left we reason by contraposition: suppose that $\theta \nvdash \bigvee \widehat{Z}$. Then by Lindenbaum's lemma there is a maximal consistent set of formulas Γ containing θ and $\neg \bigvee \widehat{Z}$. Then $\Gamma \cap \Sigma$ is an atom, and is clearly consistent with θ. But it cannot be consistent with $\bigvee \widehat{Z}$: since any two distinct atoms are mutually inconsistent, this could only be the case if in fact $\Gamma \cap \Sigma \in Z$, which implies that $\widehat{\Gamma \cap \Sigma} \vdash \bigvee \widehat{Z}$. Since $\widehat{\Gamma \cap \Sigma} \in \Gamma$ we would then get $\bigvee \widehat{Z} \in \Gamma$, and since we had $\neg \bigvee \widehat{Z} \in \Gamma$ this is a contradiction since Γ was consistent. ∎

DEFINITION 4. Given any label π, we define the relation $S_\pi^\Sigma \subseteq \mathsf{At}(\Sigma) \times \mathcal{P}(\mathsf{At}(\Sigma))$ by setting $(w, Z) \in S_\pi^\Sigma$ iff $\widehat{w} \wedge \langle \pi \rangle (\widehat{Z}; \bigvee \widehat{Z})$ is consistent with respect to the system Ax3. The *canonical neighborhood model* over Σ, denoted $\mathfrak{C},^\Sigma$ is defined as the triple $(W^\Sigma, R^\Sigma, V^\Sigma)$ where W^Σ is the set of atoms over Σ, $R_a^\Sigma = S_a^\Sigma$ for each atomic label a, and $V^\Sigma(p) = \{w \in W^\Sigma \mid p \in w\}$.

The key lemma in the completeness proof, which is proved using the induction rules for the Kleene star, is the following:

LEMMA 4. *For each label π, we have $S_{\pi*}^\Sigma \subseteq [\![*]\!](S_\pi^\Sigma)$.*

PROOF. Since the set of atoms is finite, we can use the characterization of the Kleene star operation on finite models given by Proposition 4.

Suppose that $(w, Z) \in S_{\pi*}^\Sigma$, meaning that $\nvdash \neg(\widehat{w} \wedge \langle \pi^* \rangle (\widehat{Z}; \bigvee \widehat{Z}))$. Let $\gamma[Z]$ be the disjunction of all formulas $\widehat{v}$ for $(v, Z) \in [\![*]\!](S_\pi^\Sigma)$. We want to show that $\langle \pi^* \rangle (\widehat{Z}; \bigvee \widehat{Z}) \vdash \gamma[Z]$. It will then follow that $\widehat{w} \wedge \gamma[Z]$ is consistent, and clearly since w is an atom this can only happen if $\widehat{w}$ is already a disjunct of $\gamma[Z]$ which means that $(w, Z) \in [\![*]\!](S_\pi^\Sigma)$ as desired.

More generally, for $Z' \subseteq Z$ let $\gamma[Z', Z]$ be the disjunction of all formulas $\widehat{v}$ where v is an atom such that $(v, Z'') \in [\![*]\!](S_\pi^\Sigma)$ and $Z' \subseteq Z'' \subseteq Z$ for some set Z''. We will show that $\langle \pi^* \rangle (\widehat{Z'}; \bigvee \widehat{Z}) \vdash \gamma[Z', Z]$. The special case for the formula $\gamma[Z, Z] = \gamma[Z]$ then yields the desired result.

We first prove the claim for the case of $Z' = \emptyset$. We have

$$\langle \pi^* \rangle (\widehat{\emptyset}; \bigvee \widehat{Z}) = \langle \pi^* \rangle (\emptyset; \bigvee \widehat{Z}) = \langle \pi^* \rangle \bigvee \widehat{Z}$$

So we want to show that $\langle \pi^* \rangle \bigvee \widehat{Z} \vdash \gamma[\emptyset, Z]$, and by the first induction rule it suffices to prove that $\bigvee \widehat{Z} \vdash \gamma[\emptyset, Z]$ and $\langle \pi \rangle \gamma[\emptyset, Z] \vdash \gamma[\emptyset, Z]$. Now, since $\gamma[\emptyset, Z]$ is a disjunction of conjunctions of atoms, it follows from Lemma 3

that, for any formula θ, we have that $\theta \vdash \gamma[\emptyset, Z]$ iff every atom that is consistent with θ is also consistent with $\gamma[\emptyset, Z]$.

Suppose first that w is consistent with $\bigvee \widehat{Z}$. Then w must be in Z, and since $(w, \{w\}) \in \mathsf{skip} \subseteq [\![*]\!](S_\pi^\Sigma)$, w is consistent with $\gamma[\emptyset, Z]$ as required.

Next, suppose that w consistent with $\langle \pi \rangle \gamma[\emptyset, Z]$. By Lemma 1 there must be some set Z' such that w is consistent with $\langle \pi \rangle (\widehat{Z'}; \bigvee \widehat{Z'})$ and $u \vdash \gamma[\emptyset, Z]$ for each $u \in Z'$. We get that $(w, Z') \in S_\pi^\Sigma$, and furthermore for each $u \in Z'$ there must be some $Z_u \subseteq Z$ with $(u, Z_u) \in [\![*]\!](S_\pi^\Sigma)$. We get:

$$(Z', \{Z_u \mid u \in Z'\}) \in \overline{[\![*]\!](S_\pi^\Sigma)}$$

and hence we obtain:

$$(w, \bigcup_{u \in Z'} Z_u) \in S_\pi^\Sigma [\![;]\!]([\![*]\!](S_\pi^\Sigma)) \subseteq [\![*]\!](S_\pi^\Sigma)$$

and since $\emptyset \subseteq \bigcup_{u \in Z'} Z_u \subseteq Z$, we get w consistent with $\gamma[\emptyset, Z]$ as required.

Next, consider the case where $Z' \subseteq Z$ is a singleton $\{s\}$. We write $\gamma[s, Z]$ rather than $\gamma[\{s\}, Z]$. We must show that $\langle \pi^* \rangle (\widehat{s}; \bigvee \widehat{Z}) \vdash \gamma[s, Z]$, and we use Lemma 3 as before. By the second induction rule, it suffices to prove that

$$\widehat{s} \wedge \bigvee \widehat{Z} \vdash \gamma[s, Z]$$

and

$$\langle \pi^* \rangle (\gamma[s, Z]; \langle \pi^* \rangle \bigvee \widehat{Z}) \vdash \gamma[s, Z]$$

The first statement is similar to the proof that $\bigvee \widehat{Z} \vdash \gamma[\emptyset, Z]$ so we leave it out. For the second part, suppose that the atom w is consistent with the formula $\langle \pi \rangle (\gamma[s, Z]; \langle \pi^* \rangle \bigvee \widehat{Z})$. By the previous argument (i.e. for the case where $Z' = \emptyset$) we get

$$\langle \pi^* \rangle \bigvee \widehat{Z} \vdash \gamma[\emptyset, Z]$$

so by (Mon) we find that w is consistent with $\langle \pi \rangle (\gamma[s, Z]; \gamma[\emptyset, Z])$. By Lemma 1, there is some set Y such that w is consistent with $\langle \pi \rangle (\widehat{Y}; \bigvee \widehat{Y})$, $u \vdash \gamma[\emptyset, Z]$ for each $u \in Y$, and $v \vdash \gamma[s, Z]$ for some $v \in Y$. We get that $(w, Y) \in S_\pi^\Sigma$. Also, there is a set Z_v with $s \in Z_v \subseteq Z$ and $(v, Z_v) \in [\![*]\!](S_\pi^\Sigma)$, and for each $u \neq v$ in Y there is some $Z_u \subseteq Z$ such that $(u, Z_u) \in [\![*]\!](S_\pi^\Sigma)$. If we set:

$$F = \{Z_v\} \cup \{Z_u \mid u \in Y \setminus \{v\}\}$$

then we get $\{s\} \subseteq \bigcup F \subseteq Z$. Furthermore, we get

$$(Y, F) \in \overline{[\![*]\!](S_\pi^\Sigma)}$$

and hence we obtain the required inclusion:

$$(w, \bigcup F) \in S_\pi^\Sigma [\![;]\!]([\![*]\!](S_\pi^\Sigma)) \subseteq [\![*]\!](S_\pi^\Sigma)$$

Finally, let $Z' \subseteq Z$ be an arbitrary non-empty set, and suppose w is consistent with $\langle \pi^* \rangle (\widehat{Z'}; \bigvee \widehat{Z})$, where $Z' = \{s_1, ..., s_n\}$. Then by the axiom (Fix), w is consistent with the formula

$$(\bigwedge \widehat{Z'} \wedge \bigvee \widehat{Z}) \vee \langle \pi; \pi^* \rangle (\widehat{Z'}; \bigvee \widehat{Z})$$

So it now suffices to prove that:

$$(\bigwedge \widehat{Z'} \wedge \bigvee \widehat{Z}) \vdash \gamma[Z', Z]$$

and

$$\langle \pi; \pi^* \rangle (\widehat{Z'}; \bigvee \widehat{Z}) \vdash \gamma[Z', Z]$$

Once again, the first claim follows by a familiar argument using $\mathsf{skip} \subseteq [\![*]\!](S_\pi^\Sigma)$. For the second claim, it suffices by axiom (Cmp) to prove that:

$$\langle \pi \rangle (\langle \pi^* \rangle (\widehat{s}_1; \bigvee \widehat{Z}), ..., \langle \pi^* \rangle (\widehat{s}_n; \bigvee \widehat{Z}); \langle \pi^* \rangle \bigvee \widehat{Z}) \vdash \gamma[Z', Z]$$

But, using the previous arguments together with the rule (Mon), we find that it suffices to prove:

$$\langle \pi \rangle (\gamma[s_1; Z], ..., \gamma[s_n; Z]; \gamma[\emptyset, Z]) \vdash \gamma[Z', Z]$$

We show that every atom consistent with the formula on the left-hand side is also consistent with the formula on the right-hand side. Suppose that w is consistent with the formula $\langle \pi \rangle (\gamma[s_1; Z], ..., \gamma[s_n; Z]; \gamma[\emptyset, Z])$. By Lemma 1 there must be some set Y such that w is consistent with $\langle \pi \rangle (\widehat{Y}; \bigvee \widehat{Y})$, $u \vdash \gamma[\emptyset, Z]$ for each $u \in Y$, and for each $i \in \{1, ..., n\}$ we have $v_i \vdash \gamma[s_i, Z]$ for some $v_i \in Y$. We get that $(w, Y) \in S_\pi^\Sigma$. Furthermore for each $i \in \{1, ..., n\}$ there is some set S_i such that $s_i \in S_i \subseteq Z$ and $(v_i, S_i) \in [\![*]\!](S_\pi^\Sigma)$, and for each $u \notin \{v_1, ..., v_n\}$, $u \in Y$, there is some $Z_u \subseteq Z$ such that $(u, Z_u) \in [\![*]\!](S_\pi^\Sigma)$. If we now set:

$$F = \{S_1, ..., S_n\} \cup \{Z_u \mid u \in Y \setminus \{v_1, ..., v_n\}\}$$

then we get $\{s_1, ..., s_n\} \subseteq \bigcup F \subseteq Z$. Furthermore, we get

$$(Y, F) \in \overline{[\![*]\!](S_\pi^\Sigma)}$$

and hence we obtain:

$$(w, \bigcup F) \in S_\pi^\Sigma [\![;]\!]([\![*]\!](S_\pi^\Sigma)) \subseteq [\![*]\!](S_\pi^\Sigma)$$

as required. ∎

Lemma 4 is needed to prove Lemma 5 below, by induction on the complexity of program terms. Say that a label π is *safe* if, for every formula γ such that the term γ? appears in π, we have that $\gamma \in \Sigma$ and furthermore, $\gamma \in w$ iff $\mathfrak{C}^{\Sigma}, w \Vdash \gamma$ for each $w \in \mathsf{At}(\Sigma)$.

LEMMA 5. *For every safe label π, we have $S_{\pi}^{\Sigma} \subseteq R_{\pi}^{\Sigma}$.*

PROOF. By induction on the complexity of safe labels. For γ?, the result follows from the safety assumption and the observation that

$$S_{\gamma?}^{\Sigma} = [\![?]\!]\{w \mid \mathsf{At}(\Sigma) \mid \gamma \in w\}$$

This observation can be proved as follows: since γ is safe we have $\gamma \in \Sigma$, so $\widehat{w} \wedge \langle\gamma?\rangle(\widehat{Z}, \bigvee \widehat{Z})$ is consistent iff $\widehat{w} \wedge \bigwedge \widehat{Z} \wedge \bigvee \widehat{Z}$ is consistent, iff $\gamma \in w$ and $\widehat{Z} = \{w\}$ since w is an atom and $\widehat{Z}$ a set of atoms. Hence $S_{\gamma?}^{\Sigma} = \{(w, \{w\}) \mid \gamma \in w\}$ and the result follows from the definition of $[\![?]\!]$.

For the Kleene star, by Lemma 4 we have $S_{\pi*}^{\Sigma} \subseteq [\![*]\!](S_{\pi}^{\Sigma})$ for each label π. Similarly we may prove: $S_{\pi_1 \cup \pi_2}^{\Sigma} \subseteq S_{\pi_1}^{\Sigma} [\![\cup]\!] S_{\pi_2}^{\Sigma}$ and $S_{\pi_1;\pi_2}^{\Sigma} \subseteq S_{\pi_1}^{\Sigma} [\![;]\!] S_{\pi_2}^{\Sigma}$. We omit the easy argument for $\cup$. For $\cap$, suppose that $\widehat{w} \wedge \langle\pi_1 \cap \pi_2\rangle(\widehat{Z}; \bigvee \widehat{Z})$ is consistent. Then there are sets Z_1, Z_2 such that $Z = Z_1 \cup Z_2$ such that:

$$\widehat{w} \wedge \langle\pi_1\rangle(\widehat{Z_1}; \bigvee \widehat{Z}) \wedge \langle\pi_2\rangle(\widehat{Z_2}; \bigvee \widehat{Z})$$

is consistent. Hence both $\widehat{w} \wedge \langle\pi_1\rangle(\widehat{Z_1}; \bigvee \widehat{Z})$ and $\widehat{w} \wedge \langle\pi_2\rangle(\widehat{Z_2}; \bigvee \widehat{Z})$ are consistent, and using Lemma 1 we find sets $Y_1, Y_2 \subseteq \mathsf{At}(\Sigma)$ (corresponding to disjuncts of the normal form) such that $Z_1 \subseteq Y_1 \subseteq Z$ and $Z_2 \subseteq Y_2 \subseteq Z$ and such that both $\widehat{w} \wedge \langle\pi_1\rangle(\widehat{Y_1}; \bigvee \widehat{Y_1})$ and $\widehat{w} \wedge \langle\pi_2\rangle(\widehat{Y_2}; \bigvee \widehat{Y_2})$ are consistent. Hence $(w, Y_1) \in S_{\pi_1}^{\Sigma}$ and $(w, Y_2) \in S_{\pi_2}^{\Sigma}$, hence $(w, Y_1 \cup Y_2) \in S_{\pi_1}^{\Sigma} [\![\cap]\!] S_{\pi_2}^{\Sigma}$. The result now follows since clearly $Y_1 \cup Y_2 = Z$.

For composition, suppose that atom w is consistent with the formula $\langle\pi_1;\pi_2\rangle(\widehat{Z}; \bigvee \widehat{Z})$, where $Z = \{v_1, ..., v_n\}$. Then w is consistent with

$$\langle\pi_1\rangle(\langle\pi_2\rangle(\widehat{v_1}; \bigvee \widehat{Z}), ..., \langle\pi_2\rangle(\widehat{v_n}; \bigvee \widehat{Z}); \langle\pi_2\rangle \bigvee \widehat{Z})$$

by the axiom (Cmp). For each $i \in \{1, ..., m\}$ let δ_i be the disjunction of the set of all formulas $\widehat{u}$ such that u is an atom with $(u, U) \in S_{\pi_2}^{\Sigma}$ for some set of atoms U with $v_i \in U$ and $U \subseteq Z$, and let θ be the disjunction of all formulas $\widehat{u}$ such that u is an atom with $(u, U) \in S_{\pi_2}^{\Sigma}$ for some $U \subseteq Z$.

We first claim that:

$$\langle\pi_1\rangle(\langle\pi_2\rangle(\widehat{v_1}; \bigvee \widehat{Z}), ..., \langle\pi_2\rangle(\widehat{v_n}; \bigvee \widehat{Z}); \langle\pi_2\rangle \bigvee \widehat{Z}) \vdash \langle\pi_1\rangle(\delta_1, ..., \delta_n; \theta)$$

To see this, let the maximum modal depth of formulas in Σ be k, and let F_Σ^{2+k} be the set of all formulas of modal depth at most $2 + k$, such that only labels appearing in formulas in Σ may appear in formulas in F_Σ^{2+k}. Let an *extended atom* be a maximal consistent subset of F_Σ^{2+k}. Since there are only finitely many formulas in F_Σ^{2+k} up to provable equivalence, there are at most finitely many extended atoms, and for each extended atom e we can form the conjunction $\widehat{e}$ of all formulas in e "up to logical equivalence", picking one conjunct from each logical equivalence class. Since both formulas $\langle \pi_1 \rangle (\langle \pi_2 \rangle (\widehat{v_1}; \bigvee \widehat{Z}), ..., \langle \pi_2 \rangle (\widehat{v_n}; \bigvee \widehat{Z}); \langle \pi_2 \rangle \bigvee \widehat{Z})$ and $\langle \pi_1 \rangle (\delta_1, ..., \delta_n; \theta)$ are of modal depth $\leq 2 + k$, it suffices to prove that every extended atom e which contains the following formula:

$$\langle \pi_1 \rangle (\langle \pi_2 \rangle (\widehat{v_1}; \bigvee \widehat{Z}), ..., \langle \pi_2 \rangle (\widehat{v_n}; \bigvee \widehat{Z}); \langle \pi_2 \rangle \bigvee \widehat{Z})$$

also contains:

$$\langle \pi_1 \rangle (\delta_1, ..., \delta_n; \theta).$$

So let e be an extended atom containing the first of these two formulas. Once again, a proof similar to that of Lemma 1, shows that the formula

$$\langle \pi_1 \rangle (\langle \pi_2 \rangle (\widehat{v_1}; \bigvee \widehat{Z}), ..., \langle \pi_2 \rangle (\widehat{v_n}; \bigvee \widehat{Z}); \langle \pi_2 \rangle \bigvee \widehat{Z})$$

is equivalent to a disjunction of formulas of the form $\langle \pi_1 \rangle (\widehat{E}, \bigvee \widehat{E})$ where E is a set of extended atoms such that $\langle \pi_2 \rangle \bigvee \widehat{Z} \in \bigcap E$ and $\langle \pi_2 \rangle (\widehat{v_i}; \bigvee \widehat{Z}) \in \bigcup E$ for each $i \in \{1, ..., n\}$. So one of these disjuncts $\langle \pi_1 \rangle (\widehat{E}, \bigvee \widehat{E})$ belongs to e. Furthermore, it is not hard to show that $\vdash \widehat{e'} \to \theta$ for each $e' \in E$: if $e' \in E$ then $\langle \pi_2 \rangle \bigvee \widehat{Z} \in e'$. So we show that $\langle \pi_2 \rangle \bigvee \widehat{Z} \vdash \theta$: recall that θ was the disjunction of all formulas $\widehat{u}$ such that u is an atom with $(u, U) \in S_{\pi_2}^\Sigma$ for some $U \subseteq Z$. We show that any atom u' consistent with $\langle \pi_2 \rangle \bigvee \widehat{Z}$ is consistent with θ also, from which the desired conclusion follows using Lemma 3. But if u' is consistent with $\langle \pi_2 \rangle \bigvee \widehat{Z}$ then by Lemma 1 we find a subset $U \subseteq Z$ such that u' is consistent with $\langle \pi_2 \rangle (\widehat{U}, \bigvee \widehat{U})$. Then $(u, U) \in S_{\pi_2}^\Sigma$, hence u must be consistent with θ as required.

Similarly one can show that $\vdash \widehat{e'} \to \delta_i$ for each $e' \in E$ such that $\langle \pi_2 \rangle (\widehat{v_i}; \bigvee \widehat{Z}) \in e'$ (since $e' \cap \Sigma$ is an atom consistent with $\langle \pi_2 \rangle (\widehat{v_i}; \bigvee \widehat{Z})$). Therefore, we get:

$$\langle \pi_1 \rangle (\widehat{E}, \bigvee \widehat{E}) \vdash \langle \pi_1 \rangle (\delta_1, ..., \delta_n; \theta)$$

by (Mon), whence $\langle \pi_1 \rangle (\delta_1, ..., \delta_n; \theta)$ belongs to e as well.

So w is consistent with the formula $\langle\pi_1\rangle(\delta_1, ..., \delta_n; \theta)$, and by Lemma 1 there is a set Q of atoms such that w is consistent with $\langle\pi_1\rangle(\widehat{Q}; \bigvee\widehat{Q})$, $s \vdash \theta$ for each $s \in Q$ and for each $i \in \{1, ..., n\}$ there is $t_i \in Q$ such that $t_i \vdash \delta_i$. It follows from this that for each $s \in Q$ there is some $U_s \subseteq Z$ such that $(s, U_s) \in S_{\pi_2}^\Sigma$, and for each $i \in \{1, ..., n\}$ there is some $P_i \subseteq Z$ such that $v_i \in P_i$ and $(t_i, P_i) \in S_{\pi_2}^\Sigma$. If we set

$$F = \{U_s \mid s \in Q\} \cup \{P_i \mid i \in \{1, ..., n\}\}$$

then we get $(Q, F) \in \overline{S_{\pi_2}^\Sigma}$, and so $(w, \bigcup F) \in S_{\pi_1}^\Sigma [\![;]\!] S_{\pi_2}^\Sigma$. But $\bigcup F = Z$, so we get $(w, Z) \in S_{\pi_1}^\Sigma [\![;]\!] S_{\pi_2}^\Sigma$ as required.

Finally, a straightforward induction now shows that $S_\pi^\Sigma \subseteq R_\pi^\Sigma$ for each safe label π, using monotonicity of each of the operations $[\![\cup]\!], [\![\cap]\!], [\![;]\!], [\![*]\!]$. For atomic labels the claim holds by definition of $R_a^\Sigma = S_a^\Sigma$. For the case of iteration, as an example, we have:

$$\begin{aligned} S_{\pi^*}^\Sigma &\subseteq [\![*]\!](S_\pi^\Sigma) \\ &\subseteq [\![*]\!](R_\pi^\Sigma) \\ &= R_{\pi^*}^\Sigma \end{aligned}$$

The other cases are similar.   ∎

Using Lemma 5 we can prove a truth lemma for the canonical model:

LEMMA 6. *For every atom w and any $\psi \in \Sigma$, we have the equivalence $(\mathfrak{C}^\Sigma, w) \Vdash \psi$ if and only if $\psi \in w$.*

PROOF. By induction on the complexity of ψ. Note that the induction hypothesis for subformulas of ψ guarantees that every label appearing in ψ is safe. The only interesting cases are formulas of the form $\langle\pi\rangle(\Psi; \varphi)$.

For right to left, suppose $\langle\pi\rangle(\Psi; \varphi) \in w$. By Lemma 1 we find a set Z of atoms such that $\langle\pi\rangle(\widehat{Z}, \bigvee\widehat{Z})$ is consistent with w, hence $(w, Z) \in S_\pi^\Sigma$, and such that $\Psi \subseteq \bigcup Z$ and $\varphi \in \bigcap Z$. By Lemma 5 we get $(w, Z) \in R_\pi^\Sigma$, and the induction hypothesis applied to the formulas in $\Psi \cup \{\varphi\}$ now readily yields $\mathfrak{C}^\Sigma, w \Vdash \langle\pi\rangle(\Psi; \varphi)$ as required.

For left to right, it suffices to show that for all formulas $\langle\pi\rangle(\Psi; \varphi) \in \Sigma$, all sets of atoms Z and all atoms w such that $(w, Z) \in R_\pi^\Sigma$, $\varphi \in \bigcap Z$ and $\Psi \subseteq \bigcup Z$, we have $\langle\pi\rangle(\Psi; \varphi) \in w$. The required result then follows by applying the induction hypothesis to Ψ, φ. We prove the claim by induction on the complexity of the label π, assuming that π is a safe label.

If π is an atomic label a then we have $R_a^\Sigma = S_a^\Sigma$. So if $(w, Z) \in R_a^\Sigma$ then $(w, Z) \in S_a^\Sigma$, so w is consistent with $\langle a\rangle(\widehat{Z}; \bigvee\widehat{Z})$. From this we can easily

derive that w is consistent with $\langle a \rangle (\Psi; \varphi)$ by an argument combining the rule (Mon) and the axiom (Weak), given that $\varphi \in \bigcap Z$ and $\Psi \subseteq \bigcup Z$. Since $\langle a \rangle (\Psi; \varphi) \in \Sigma$ and w is an atom it follows that $\langle a \rangle (\Psi; \varphi) \in w$ as required.

The induction steps for the constructions of test, choice, parallel composition and sequential composition are straightforward, making use of Fischer-Ladner closure of Σ at each step.

We now focus on the case of the Kleene star. Suppose that there is some Z such that $(w, Z) \in R^{\Sigma}_{\pi^*}$, $\varphi \in \bigcap Z$ and $\Psi \subseteq \bigcup Z$. By Proposition 4 there is some natural number n with $(w, Z) \in (R^{\Sigma}_\pi)^{[n]}$, so we reason by induction on n. That is, we show that for all w, Z, Ψ, φ and all $n \in \omega$, if $(w, Z) \in (R^{\Sigma}_\pi)^{[n]}$, $\varphi \in \bigcap Z$ and $\Psi \subseteq \bigcup Z$, then $\langle \pi^* \rangle (\Psi; \varphi) \in w$.

For $n = 0$, the result holds trivially, as $(R^{\Sigma}_\pi)^{[n]} = \emptyset$. And if the induction hypothesis holds for n, if $(w, Z) \in (R^{\Sigma}_\pi)^{[n+1]}$ then either $(w, Z) \in \mathsf{skip}$, or:

$$(w, Z) \in R^{\Sigma}_\pi [\![;]\!] (R^{\Sigma}_\pi)^{[n]}$$

In the first case, we have $Z = \{w\}$ so it immediately follows (using Fisher-Ladner closure of Σ and $\langle \pi^* \rangle (\Psi; \varphi) \in \Sigma$) that $\bigwedge \Psi \wedge \varphi \in w$. By the axiom (Fix), we get $\langle \pi^* \rangle (\Psi; \varphi) \in w$ as required. Otherwise, if $(w, Z) \in R^{\Sigma}_\pi [\![;]\!] (R^{\Sigma}_\pi)^{[n]}$, there is a set Y and a family of sets F with $(w, Y) \in R^{\Sigma}_\pi$, $(Y, F) \in \overline{(R^{\Sigma}_\pi)^{[n]}}$, $\bigcup F = Z$, $\varphi \in \bigcap X$ for each $X \in F$, and for each $\psi \in \Psi$ there is some $X_\psi \in F$ with $\psi \in \bigcup X_\psi$. By Fisher-Ladner closure we get

$$(\bigwedge \Psi \wedge \varphi) \vee \langle \pi; \pi^* \rangle (\Psi; \varphi) \in \Sigma$$

and hence, with $\Psi = \{\psi_1, ..., \psi_k\}$:

$$\langle \pi \rangle (\langle \pi^* \rangle (\psi_1; \varphi), ..., \langle \pi^* \rangle (\psi_k; \varphi); \langle \pi^* \rangle \varphi) \in \Sigma$$

For every $\psi \in \Psi$ and every $v \in Y$ such that $(v, X_\psi) \in (R^{\Sigma}_\pi)^{[n]}$, since we have $\psi \in \bigcup X_\psi$ and $\varphi \in \bigcap X_\psi$ we obtain from the "inner" induction hypothesis on n that $\langle \pi^* \rangle (\psi; \varphi) \in v$. By a similar argument, we get $\langle \pi^* \rangle \varphi \in v$ for all $X \in F$ and all $v \in Y$ with $(v, X) \in (R^{\Sigma}_\pi)^{[n]}$. Since $(Y, F) \in \overline{(R^{\Sigma}_\pi)^{[n]}}$, we find that for each $\psi \in \Psi$ we have $\langle \pi^* \rangle (\psi; \varphi) \in \bigcup Y$, and we also get $\langle \pi^* \rangle \varphi \in \bigcap Y$. Since $(w, Y) \in R^{\Sigma}_\pi$ we can now apply the "outer" induction hypothesis on labels to the label π and get:

$$\langle \pi \rangle (\langle \pi^* \rangle (\psi_1; \varphi), ..., \langle \pi^* \rangle (\psi_n; \varphi); \langle \pi^* \rangle \varphi) \in w$$

By applying the axiom (Cmp) we get $\langle \pi; \pi^* \rangle (\Psi; \varphi) \in w$, and hence

$$(\bigwedge \Psi \wedge \varphi) \vee \langle \pi; \pi^* \rangle (\Psi; \varphi) \in w$$

just from classical propositional logic. Finally, by the fixpoint axiom (Fix) we get $\langle \pi^* \rangle (\Psi; \varphi) \in w$ as required. ∎

PROOF OF THEOREM 5. Suppose the formula φ is not provable, so that $\neg\varphi$ is consistent. By Lemma 2, $\neg\varphi$ belongs to some finite Fischer-Ladner closed set Σ and since $\neg\varphi$ is consistent it belongs to some atom w. Hence $\varphi \notin w$ and by Lemma 6 we have $\mathfrak{C}^\Sigma, w \nVdash \varphi$. So φ is not valid. ∎

We note that as a corollary to the completeness proof, which produces a finite model of effectively bounded size for a consistent formula, we get:

THEOREM 6. IPDL *has the finite model property and is decidable.*

5.   Comparison with game logic

We now show that IPDL can, in a precise sense, be viewed as a language extension of dual-free game logic, called GL here for short. Formally, formulas of GL and game terms are defined by the following dual grammar:

$$\varphi := p \in \mathsf{Prop} \mid \varphi \wedge \varphi \mid \neg\varphi \mid \langle \pi \rangle \varphi$$

$$\pi := a \in \mathcal{A} \mid \varphi? \mid \pi;\pi \mid \pi \cup \pi \mid \pi \cap \pi \mid \pi^*$$

with Prop a fixed set of propositional variables and $\mathcal{A}$ a set of atomic games, both assumed to be countably infinite. Note that GL is a syntactic fragment of IPDL. Here, $\cup$ is interpreted as "angelic choice" (choice for Player I), $\cap$ is interpreted as "demonic choice" (choice for Player II), ; is sequential game composition and $*$ is finite game iteration (controlled by Player I).

Semantics of game logic formulas are given by neighborhood frames, with the extra constraint that the family of neighborhoods associated with a world is upwards closed under set inclusion:

DEFINITION 5. A neighborhood frame (W, R) is said to be a *monotonic power frame* if the following condition holds for each $a \in \mathcal{A}$:

(Monotonicity) For all $u \in W$, if $(u, Z) \in R_a$ and $Z \subseteq Z'$ then $(u, Z') \in R_a$.

A monotonic power model is a neighborhood model whose underlying frame is a monotonic power frame.

To provide a semantics for formulas in a model, we need to interpret the game constructors. In what follows, we shall use double vertical lines $\|-\|$ for semantic denotations of formulas in GL and game constructors in

monotonic neighborhood models, as distinct from our semantics for PDL presented earlier, that used square denotation brackets $[\![-]\!]$.

More precisely, using the format introduced in [3], we define operations on the lattice $\mathcal{N}W = \mathcal{P}(W \times \mathcal{P}(W))$ of *neighborhood relations* over W:

- $R\|\cup\|R' = R \cup R'$

- $R\|\cap\|R' = R \cap R'$

- $(u, Z') \in R\|;\|R'$ iff there is some $Z \subseteq W$ with $(u, Z) \in R$ and $(v, Z') \in R'$ for all $v \in Z$.

- $\|?\|(Z) = \{(w, Z') \in W \times \mathcal{P}(W) \mid w \in Z \cap Z'\}$

Finally, we define $\|*\|R$ to be the least fixpoint in the lattice $\mathcal{N}W$ of the monotone map F defined by:

$$FS = \mathsf{skip}^{\uparrow}\|\cup\|(R\|;\|S)$$

where $\mathsf{skip}^{\uparrow} = \{(w, Z) \in W \times \mathcal{P}(W) \mid w \in Z\}$. We can now set up the semantics of GL. Fixing a monotonic power model $\mathfrak{M}$, we define the interpretation of every formula φ and the neighborhood relations R_π corresponding to each game term π in the obvious way, so that in particular we have $R_{\pi_1 \cup \pi_2} = R_{\pi_1}\|\cup\|R_{\pi_2}$, $R_{\pi_1 \cap \pi_2} = R_{\pi_1}\|\cap\|R_{\pi_2}$ etc., and $u \in \|\langle\pi\rangle\varphi\|$ iff $(u, \|\varphi\|) \in R_a$. For a monotonic power model $\mathfrak{M} = (W, R, V)$ and $u \in W$ we shall also write $\mathfrak{M}, u \vDash \varphi$ for $u \in \|\varphi\|$. Since semantic interpretations are always defined relative to a model, if necessary we shall use the notation $\|-\|_{\mathfrak{M}}$ rather than $\|-\|$ to make it clear which model $\mathfrak{M}$ is being referred to. We write $\vDash \varphi$ if $\mathfrak{M}, u \vDash \varphi$ for every pointed monotone power model $(\mathfrak{M}, u)$.

We can now state precisely how IPDL extends dual-free game logic:

THEOREM 7. IPDL *is a conservative extension of* GL. *That is, for every* GL-*formula* φ, *we have that*

$$\vDash \varphi \ \textit{iff} \ \Vdash \varphi$$

PROOF. For every neighborhood model $\mathfrak{M}$, we define a monotonic power model $\mathfrak{M}^{\uparrow}$ as follows: let $\mathfrak{M} = (W, R, V)$. We define the monotonic power model $\mathfrak{M}^{\uparrow} = (W, \underline{R}, V)$ as follows: set $(u, Z) \in \underline{R}_a$ iff there is some $Z' \subseteq Z$ with $(u, Z') \in R_a$. Now we have the following:

CLAIM 2. For all GL-formula φ and neighborhood models $\mathfrak{M}$, $\|\varphi\|_{\mathfrak{M}^{\uparrow}} = [\![\varphi]\!]_{\mathfrak{M}}$.

PROOF OF CLAIM 2. The result follows easily once we have established:

CLAIM 3. Given a neighborhood model $\mathfrak{M}$ and a term π, let R_π denote the neighborhood relation corresponding to π in $\mathfrak{M}$ computed by applying the operations $[\![?]\!], [\![\cup]\!], [\![\cap]\!], [\![;]\!], [\![*]\!]$, and let S_π denote the neighborhood relation corresponding to π in $\mathfrak{M}^\uparrow$ computed by applying the operations $\|?\|, \|\cup\|, \|\cap\|, \|;\|, \|*\|$. Then for all w, Z, we have $(w, Z) \in S_\pi$ iff there is some $Z' \subseteq Z$ with $(w, Z') \in R_\pi$.

We devote the rest of the proof to establishing this claim. The claim is immediate for atomic games, and the step for the test operator follows trivially from the definitions. The direction from right to left is easy in each case, so we focus on the converse implication.

The induction step for $\cup$ is entirely straightforward. For $\cap$, if $(w, Z) \in S_{\pi_1 \cap \pi_2}$ then $(w, Z) \in S_{\pi_1}$ and $(w, Z) \in S_{\pi_2}$. By the induction hypothesis, there are sets $Y, Y' \subseteq Z$ such that $(w, Y) \in R_{\pi_1}$ and $(w, Y') \in R_{\pi_2}$. So $(w, Y \cup Y') \in R_{\pi_1 \cap \pi_2}$. Since $Y \cup Y' \subseteq Z$, we are done.

For the sequential composition operator, suppose $(w, Z) \in S_{\pi_1 ; \pi_2}$. Then there exists a set Y such that $(w, Y) \in S_{\pi_1}$ and $(v, Z) \in S_{\pi_2}$ for each $v \in Y$. By the induction hypothesis, there are sets $\{Z'_v\}_{v \in Y}$ with $(v, Z'_v) \in R_{\pi_2}$ and $Z'_v \subseteq Z$, and there exists a set $Y' \subseteq Y$ with $(w, Y') \in R_{\pi_1}$. We then get $(w, \bigcup\{Z'_v \mid v \in Y'\}) \in R_{\pi_1 ; \pi_2}$, and because of the successive inclusions $\bigcup\{Z'_v \mid v \in Y'\} \subseteq \bigcup\{Z'_v \mid v \in Y\} \subseteq Z$, we are done.

Finally, we consider the case of game iteration. First, we recall that skip denotes the neighborhood relation $\{(w, \{w\}) \mid w \in W\}$, and $\mathsf{skip}^\uparrow$ denotes the relation $\{(w, Z) \mid w \in Z\}$.

Suppose the induction hypothesis holds for R_π. Let

$$F := \lambda Z.\mathsf{skip}^\uparrow \|\cup\| (S_\pi \|;\| Z)$$

so that S_{π^*} is equal to the least fixpoint for F. Alternatively, we can describe S_{π^*} as the least fixpoint of the map F restricted to the complete sub-lattice of $\mathcal{N}W$ given by $\{R \in \mathcal{N}W \mid \mathsf{skip}^\uparrow \subseteq R\}$. Its bottom element is $\mathsf{skip}^\uparrow$, so we can write the approximating sequence for the least fixpoint as:

$$\mathsf{skip}^\uparrow \subseteq F\mathsf{skip}^\uparrow \subseteq F^2\mathsf{skip}^\uparrow \subseteq F^3\mathsf{skip}^\uparrow \dots \quad F^\omega\mathsf{skip}^\uparrow \subseteq F^{\omega+1}\mathsf{skip}^\uparrow \dots$$

We denote the first two entries in the series as $F^0\mathsf{skip}^\uparrow$ and $F^1\mathsf{skip}^\uparrow$. We show, by transfinite induction, that

$(w, Z) \in F^\xi\mathsf{skip}^\uparrow$ iff there is some $Z' \subseteq Z$ such that $(w, Z') \in (R_\pi)^{[\xi]}$

The result then follows by considering ξ such that $R_{\pi^*} = (R_\pi)^{[\xi]}$ and γ such that $S_{\pi^*} = F^\gamma \mathsf{skip}^\uparrow$. Next, pick some ρ greater than both γ and ξ. Then we are done, because we have $R_{\pi^*} = R_\pi^{[\rho]}$ and $S_{\pi^*} = F^\rho \mathsf{skip}^\uparrow$.

To establish the claim, the case for $\xi = 0$ is trivial since by definition $F^0 \mathsf{skip}^\uparrow = \mathsf{skip}^\uparrow$ and $R_\pi^{[0]} = \mathsf{skip}$. Successor ordinals $\xi + 1$ are handled by unfolding and comparing the definitions of $R_\pi^{[\xi+1]}$ and $F^{\xi+1} \mathsf{skip}^\uparrow$, applying the "inner" induction hypothesis to $F^\xi \mathsf{skip}^\uparrow$, applying the "outer" induction hypothesis to S_π, and then repeating and combining the previous arguments for $\cup$ and ;. Finally, limit ordinals κ are handled by simply noting that $R_\pi^{[\kappa]} = \bigcup_{\xi < \kappa} (R_\pi)^{[\xi]}$ and $F^\kappa \mathsf{skip}^\uparrow = \bigcup_{\xi < \kappa} F^\xi \mathsf{skip}^\uparrow$. ∎

We can now prove Theorem 7. Suppose φ is a formula of GL and $\Vdash \varphi$. Then since every monotonic power frame is a neighborhood frame, we get $\mathfrak{M}, w \Vdash \varphi$ for every pointed monotonic power model $(\mathfrak{M}, w)$. But if $\mathfrak{M}$ is a monotonic power model, we have $\mathfrak{M}^\uparrow = \mathfrak{M}$, so it follows from Claim 2 that $\mathfrak{M}, w \vDash \varphi$ for every pointed monotonic power model as well. Hence $\vDash \varphi$.

Conversely, suppose $\vDash \varphi$, so that φ is valid on every monotonic power frame. Then for any neighborhood model $\mathfrak{M}$ and every state w in W, we have $\mathfrak{M}^\uparrow, w \vDash \varphi$, so $\mathfrak{M}, w \Vdash \varphi$ by Claim 2. Hence $\Vdash \varphi$ as required. ∎

In other words: the formulas of IPDL that are valid on arbitrary neighborhood frames form a conservative extension of the GL-formulas that are valid over monotonic power frames.

6. Concluding remarks

In this paper, we have introduced a new propositional dynamic logic IPDL defined over instantial neighborhood logic, as a tool for exploring a new perspective on open systems computation. We found program operations that respect a natural notion of bisimulation in this setting, and we axiomatized the complete logic, which presented some non-trivial and interesting deviations from the usual proof format for PDL. Finally, we positioned our logic with respect to related views of computation by completely clarifying its relation to Parikh's dual-free game logic.

Our system fits in a broader technical context. Various extensions of our base language would make sense, notably, the addition of least and greatest fixpoint operators. Just as standard PDL can be translated into the modal μ-calculus, our logic IPDL can be translated into the extension of INL with fixpoints, a translation that is implicit in our axiom system for IPDL. The fixpoint extension of INL is very well behaved from a co-algebraic perspective.

As shown in [6], INL is a coalgebraic modal logic corresponding to a weak pullback preserving functor - the double covariant powerset functor - that additionally preserves finite sets. This means that the μ-calculus extension of INL inherits a number of properties that hold in much wider generality. In particular, it has the finite model property and it is decidable [22], and a sound and complete system of axioms is available [12]. However, as usual, such general results need not transfer to natural fragments that zoom in more closely on computation. Examples are Reynold's highly non-trivial completeness proof for CTL* [21], or Parikh's game logic, which still lacks a complete system of axioms. A closer comparison for our system would be coalgebraic PDL, [18], but there the coalgebraic type functor is a monad. By contrast, this is not the case for INL[1]. Still, there is work to be done. For instance, our sequential program composition resembles the standard Kleisli composition for the powerset functor – we leave these issues to future investigation.

These are not the only connections to be clarified. In follow-up work, we intend to show that IPDL can also throw new light on other logical systems for computation, such as concurrent PDL ([20, 7, 15]), and that it can contribute to a more fine-structured analysis of game equivalence and powers of players, linking up with game theory (see [4], for which an extended follow-up manuscript is in preparation). The relationship between IPDL and the alternating-time temporal logics ATL and ATL* also remains to be explored. Standard ATL describes abilities of players to force conditions on (potentially infinite) computations by some strategy. An INL-like extension of ATL could then allow reasoning involving more complex quantification over infinite computations, involving both universal and existential quantification over the set of computations that are compatible with a given strategy.

References

[1] L. Aceto, A. Ingólfsdóttir, K. G. Larsen, and J. Srba. *Reactive systems: modelling, specification and verification.* Cambridge University Press, 2007.

[2] R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *Journal of the ACM (JACM)*, 49(5):672–713, 2002.

[3] J. van Benthem. *Logic in games.* MIT Press, Cambridge, MA, 2014.

[4] J. van Benthem, N. Bezhanishvili, and S. Enqvist. A new game equivalence and its modal logic. In *Proceedings Sixteenth Conference on Theoretical Aspects of Rationality and Knowledge, TARK 2017, Liverpool, UK, 24-26 July 2017.*, pages 57–74, 2017.

---

[1]We came to this conclusion after a private conversation with Helle Hansen, who correctly guessed that the double covariant powerset functor is not a monad. This fact has very recently been verified by Klin and Salamanca in [16].

[5] J. van Benthem, N. Bezhanishvili, and S. Enqvist. A propositional dynamic logic for instantial neighborhood models. In A. Baltag, J. Seligman, and T. Yamada, editors, *Logic, Rationality, and Interaction, LORI 2017, Proceedings*, volume 10455 of *Lecture Notes in Computer Science*, pages 137–150. Springer, 2017.

[6] J. van Benthem, N. Bezhanishvili, S. Enqvist, and J. Yu. Instantial neighborhood logic. *The Review of Symbolic Logic*, 10(1):116–144, 2017.

[7] J. van Benthem, S. Ghosh, and F. Liu. Modelling simultaneous games in dynamic logic. *Synthese*, 165(2):247–268, 2008.

[8] D. Berwanger. Game logic is strong enough for parity games. *Studia Logica*, 75(2):205–219, 2003.

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

[10] F. Carreiro and Y. Venema. PDL inside the μ-calculus: A syntactic and an automata-theoretic characterization. *Advances in Modal Logic*, 10:74–93, 2014.

[11] C. Cîrstea, C. Kupke, and D. Pattinson. EXPTIME tableaux for the coalgebraic μ-calculus. In E. Grädel and R. Kahle, editors, *Computer Science Logic (CSL 2009)*, volume 5771 of *Lecture Notes in Computer Science*, pages 179–193. Springer, 2009.

[12] S. Enqvist, F. Seifan, and Y. Venema. Completeness for coalgebraic fixpoint logic. In *Proceedings of the 25th EACSL Annual Conference on Computer Science Logic (CSL 2016)*, volume 62 of *LIPIcs*, pages 7:1–7:19, 2016.

[13] S. Enqvist, F. Seifan, and Y. Venema. Completeness for μ-calculi: a coalgebraic approach. Technical Report PP-2017-04, Institute for Logic, Language and Computation, Universiteit van Amsterdam, 2017.

[14] G. Fontaine, R. Leal, and Y. Venema. Automata for coalgebras: An approach using predicate liftings. In *Automata, Languages and Programming: 37th International Colloquium ICALP'10*, volume 6199 of *LNCS*, pages 381–392. Springer, 2010.

[15] D. Harel, D. Kozen, and J. Tiuryn. *Dynamic logic*. MIT press, 2000.

[16] B. Klin and J. Salamanca. Iterated covariant powerset is not a monad. Available in pre-proceedings of Mathematical Foundations of Programming Semantics (MFPS 2018) at: https://www.mathstat.dal.ca/mfps2018/preproc/.

[17] D. Kozen. Results on the propositional μ-calculus. *Theoretical Computer Science*, 27:333–354, 1983.

[18] C. Kupke and H. H. Hansen. Weak completeness of coalgebraic dynamic logics. In *FICS 2015: Proceedings of the 10th International Workshop on Fixed Points in Computer Science, Berlin, Germany, 11-12 September 2015*. Cornell university Library, 2015.

[19] R. Parikh. The logic of games and its applications. *Annals of Discrete Mathematics*, 24:111–139, 1985.

[20] D. Peleg. Concurrent dynamic logic. *Journal of the ACM (JACM)*, 34(2):450–479, 1987.

[21] M. Reynolds. An axiomatization of full computation tree logic. *Journal of Symbolic Logic*, pages 1011–1057, 2001.

[22] Y. Venema. Automata and fixed point logic: a coalgebraic perspective. *Information and Computation*, 204:637–678, 2006.