Intuitionistic μ-calculus with the Lewis arrow

LLAMA seminar, University of Amsterdam

Lide Grotenhuis, joint work in progress with Bahareh Afshari

1 May 2024

ILLC, Universiteit van Amsterdam

Contents

Introduction

Fixpoint modal logics contain formulas that satisfy an equivalence of the form

$$\psi \equiv \varphi(\psi).$$

Fixpoint modal logics contain formulas that satisfy an equivalence of the form

$$\psi \equiv \varphi(\psi).$$

Example. In temporal logic, the formula $A \, \mathsf{U} \, B$ ("A until B") satisfies $A \, \mathsf{U} \, B \equiv B \vee (A \wedge \mathsf{X}(A \, \mathsf{U} \, B))$.

Fixpoint modal logics contain formulas that satisfy an equivalence of the form

$$\psi \equiv \varphi(\psi).$$

Example. In temporal logic, the formula $A \,\mathsf{U}\, B$ ("A until B") satisfies $A \,\mathsf{U}\, B \equiv B \vee (A \wedge \mathsf{X}(A \,\mathsf{U}\, B))$.

The modal μ-calculus extends basic modal logic by *fixpoint operators* μ and ν:

- $\mu X.\varphi(X)$ denotes the least fixpoint of $\varphi(X)$,
- $\nu X.\varphi(X)$ denotes the greatest fixpoint of $\varphi(X)$.

Fixpoint modal logics contain formulas that satisfy an equivalence of the form

$$\psi \equiv \varphi(\psi).$$

Example. In temporal logic, the formula $A \, \mathsf{U} \, B$ ("A until B") satisfies $A \, \mathsf{U} \, B \equiv B \vee (A \wedge \mathsf{X}(A \, \mathsf{U} \, B))$.

The modal μ-calculus extends basic modal logic by *fixpoint operators* μ and ν:

- $\mu X.\varphi(X)$ denotes the least fixpoint of $\varphi(X)$,
- $\nu X.\varphi(X)$ denotes the greatest fixpoint of $\varphi(X)$.

The logic is very expressive with many nice properties such as: decidability, the small model property, a finitary axiomatization, and uniform interpolation.

Fixpoint modal logics contain formulas that satisfy an equivalence of the form

$$\psi \equiv \varphi(\psi).$$

Example. In temporal logic, the formula $A \,\mathsf{U}\, B$ (“A until B”) satisfies $A \,\mathsf{U}\, B \equiv B \vee (A \wedge \mathsf{X}(A \,\mathsf{U}\, B))$.

The modal μ-calculus extends basic modal logic by *fixpoint operators* μ and ν:

- $\mu X.\varphi(X)$ denotes the least fixpoint of $\varphi(X)$,
- $\nu X.\varphi(X)$ denotes the greatest fixpoint of $\varphi(X)$.

The logic is very expressive with many nice properties such as: decidability, the small model property, a finitary axiomatization, and uniform interpolation.

Non-wellfounded and cyclic proof systems provide natural syntactic characterisations of the modal μ-calculus and its fragments.

Introduction: intuitionistic fixpoints

Recently, modal fixpoints over an intuitionistic propositional base have gained attention:

- Intuitionistic linear-time temporal logic (Balbiani, Boudou, Diégues & Fernández-Duque, 2019, 2022);
- Intuitionistic common-knowledge logic (Jäger & Marti, 2016);
- Intuitionistic Gödel-Löb logic (Das, van der Giessen, Marin, 2023).

Introduction: intuitionistic fixpoints

Recently, modal fixpoints over an intuitionistic propositional base have gained attention:

- Intuitionistic linear-time temporal logic (Balbiani, Boudou, Diégues & Fernández-Duque, 2019, 2022);
- Intuitionistic common-knowledge logic (Jäger & Marti, 2016);
- Intuitionistic Gödel-Löb logic (Das, van der Giessen, Marin, 2023).

Our aim. Develop a general framework and proof-theoretic techniques for studying intuitionistic modal fixpoint logics.

In the past, we have studied proof systems for:

1. intuitionistic linear-time temporal logic (Afshari, G., Leigh & Zenger, 2023);
2. intuitionistic modal logic with the master modality (Afshari, G., Leigh & Zenger, 2024, preprint).

Introduction: intuitionistic fixpoints

Recently, modal fixpoints over an intuitionistic propositional base have gained attention:

- Intuitionistic linear-time temporal logic (Balbiani, Boudou, Diégues & Fernández-Duque, 2019, 2022);
- Intuitionistic common-knowledge logic (Jäger & Marti, 2016);
- Intuitionistic Gödel-Löb logic (Das, van der Giessen, Marin, 2023).

Our aim. Develop a general framework and proof-theoretic techniques for studying intuitionistic modal fixpoint logics.

In the past, we have studied proof systems for:

1. intuitionistic linear-time temporal logic (Afshari, G., Leigh & Zenger, 2023);
2. intuitionistic modal logic with the master modality (Afshari, G., Leigh & Zenger, 2024, preprint).

Current work. We study an intuitionistic version of the modal μ-calculus with a generalisation of the modal $\Box$, namely the Lewis arrow. We provide game semantics for the logic and a non-wellfounded analytic proof system.

Introduction: the Lewis arrow

Dissatisfied with material implication, Lewis (1912,1914,1918,1932) introduced several axiom systems (S1-S5) meant to formalize strict implication:

$\varphi \prec\!\!\!\!\!\!\!\!\!\!\dashv \psi$ "ψ can be inferred from φ."

Dissatisfied with material implication, Lewis (1912,1914,1918,1932) introduced several axiom systems (S1-S5) meant to formalize strict implication:

$$\varphi \prec \psi \qquad \text{“}\psi \text{ can be inferred from } \varphi.\text{”}$$

For the modern modal logician:

$$\varphi \prec \psi := \Box(\varphi \to \psi).$$

So $\Box\varphi \equiv \top \prec \varphi$.

Dissatisfied with material implication, Lewis (1912,1914,1918,1932) introduced several axiom systems (S1-S5) meant to formalize strict implication:

$\varphi \strictif \psi$ "ψ can be inferred from φ."

For the modern modal logician:

$$\varphi \strictif \psi := \Box(\varphi \to \psi).$$

So $\Box\varphi \equiv \top \strictif \varphi$.

In an intuitionistic setting, $\strictif$ is not interdefinable with $\Box$.

Dissatisfied with material implication, Lewis (1912,1914,1918,1932) introduced several axiom systems (S1-S5) meant to formalize strict implication:

$\varphi \prec \psi$ "ψ can be inferred from φ."

For the modern modal logician:

$$\varphi \prec \psi := \Box(\varphi \to \psi).$$

So $\Box\varphi \equiv \top \prec \varphi$.

In an intuitionistic setting, $\prec$ is not interdefinable with $\Box$.

- In intuitionistic provability and preservativity logic (see e.g. Iemhof 2003 and Litak & Visser 2017): given a theory T,

$$A \prec_T B \text{ iff for all } \Sigma^0_1\text{-sentences } S, \text{ if } T \vdash S \to A \text{ then } T \vdash S \to B.$$

The logic IL_μ

Syntax of IL_μ

Fix some set Prop of propositions/variables. Formulas of IL_μ are given by the grammar:

$$\varphi, \psi ::= \bot \mid \top \mid P \mid \varphi \wedge \psi \mid \varphi \vee \psi \mid \varphi \to \psi \mid \varphi \prec \psi \mid \mu X.\varphi \mid \nu X.\varphi$$

with $P, X \in \mathrm{Prop}$ and X positive in φ. We define $\Box\varphi := \top \prec \varphi$.

Note: we add both μ and ν, as these operators are not interdefinable in the intuitionistic case.

Fix some set Prop of propositions/variables. Formulas of IL_μ are given by the grammar:

$$\varphi, \psi ::= \bot \mid \top \mid P \mid \varphi \wedge \psi \mid \varphi \vee \psi \mid \varphi \to \psi \mid \varphi \dashv \psi \mid \mu X.\varphi \mid \nu X.\varphi$$

with $P, X \in \mathrm{Prop}$ and X positive in φ. We define $\Box\varphi := \top \dashv \varphi$.

Note: we add both μ and ν, as these operators are not interdefinable in the intuitionistic case.

We consider formulas φ that are clean:

1. the free and bound variables are disjoint;
2. for each bound variable X there is a unique subformula $\sigma_X X.\delta_X$ of φ.

(Algebraic) semantics of IL_μ: bi-relational models

Formulas are evaluated in bi-relational Kripke models $M = (W, \leq, R, V)$, where

1. $\leq$ is a partial order (*the intuitionistic relation*),
2. $R \subseteq W^2$ (*the modal relation*),
3. if $w \leq vRu$ then wRu (*triangle confluence*).

Formulas are evaluated in bi-relational Kripke models $M = (W, \leq, R, V)$, where

1. $\leq$ is a partial order (*the intuitionistic relation*),
2. $R \subseteq W^2$ (*the modal relation*),
3. if $w \leq vRu$ then wRu (*triangle confluence*).

The truth relation for $\to$, $\dashv$ and the fixpoint operators is defined by

$$
\begin{array}{lll}
M, s \models \varphi \to \psi & \text{iff} & \text{for all } t \geq s \text{ if } M, t \models \varphi, \text{ then } M, t \models \psi, \\
M, s \models \varphi \dashv \psi & \text{iff} & \text{for all } sRt \text{ if } M, t \models \varphi, \text{ then } M, t \models \psi, \\
M, s \models \mu X.\varphi & \text{iff} & s \in LFP(\varphi_X^M), \\
M, s \models \nu X.\varphi & \text{iff} & s \in GFP(\varphi_X^M),
\end{array}
$$

where $\varphi_X^M : \mathcal{P}(W) \to \mathcal{P}(W)$ is the function given by $S \mapsto [\![\varphi]\!]^M_{V[X \mapsto S]}$.

(Algebraic) semantics of IL_μ: bi-relational models

Formulas are evaluated in bi-relational Kripke models $M = (W, \leq, R, V)$, where

1. $\leq$ is a partial order (*the intuitionistic relation*),
2. $R \subseteq W^2$ (*the modal relation*),
3. if $w \leq vRu$ then wRu (*triangle confluence*).

The truth relation for $\to$, $\multimap$ and the fixpoint operators is defined by

$$
\begin{array}{lll}
M, s \models \varphi \to \psi & \text{iff} & \text{for all } t \geq s \text{ if } M, t \models \varphi, \text{ then } M, t \models \psi, \\
M, s \models \varphi \multimap \psi & \text{iff} & \text{for all } sRt \text{ if } M, t \models \varphi, \text{ then } M, t \models \psi, \\
M, s \models \mu X.\varphi & \text{iff} & s \in LFP(\varphi_X^M), \\
M, s \models \nu X.\varphi & \text{iff} & s \in GFP(\varphi_X^M),
\end{array}
$$

where $\varphi_X^M : \mathcal{P}(W) \to \mathcal{P}(W)$ is the function given by $S \mapsto [\![\varphi]\!]^M_{V[X \mapsto S]}$.

Crucially, we have

$$\sigma X.\varphi \equiv \varphi[\sigma X.\varphi / X]$$

for $\sigma \in \{\mu, \nu\}$.

(Algebraic) semantics of IL_μ: monotonicity and confluence properties

A key property of intuitionistic Kripke semantics is monotonicity: if $v \geq w$ and $w \models \varphi$, then $v \models \varphi$.

(Algebraic) semantics of IL_μ: monotonicity and confluence properties

A key property of intuitionistic Kripke semantics is monotonicity: if $v \geq w$ and $w \models \varphi$, then $v \models \varphi$.

To obtain monotonicity for IL_μ, models need to satisfy triangle confluence: if $w \leq vRu$ then wRu.

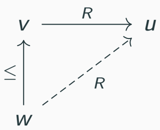

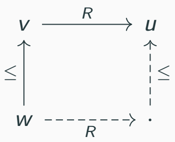

(Algebraic) semantics of IL_μ: monotonicity and confluence properties

A key property of intuitionistic Kripke semantics is monotonicity: if $v \geq w$ and $w \models \varphi$, then $v \models \varphi$.

To obtain monotonicity for IL_μ, models need to satisfy triangle confluence: if $w \leq vRu$ then wRu.

For $\Box$-formulas, forth-down confluence suffices: if $w \leq vRu$ then $wRs \leq u$ for some s.

(Algebraic) semantics of IL_μ: monotonicity and confluence properties

A key property of intuitionistic Kripke semantics is monotonicity: if $v \geq w$ and $w \models \varphi$, then $v \models \varphi$.

To obtain monotonicity for IL_μ, models need to satisfy triangle confluence: if $w \leq vRu$ then wRu.

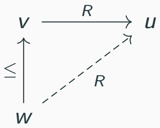

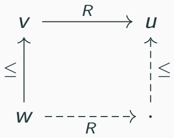

For $\Box$-formulas, forth-down confluence suffices: if $w \leq vRu$ then $wRs \leq u$ for some s.

Lemma

A $\Box$-formula φ is valid on all forth-down confluent models iff it is valid on all triangle confluent models.

Proof. Any forth-down confluent model $M = (W, \leq, R, V)$ induces a triangle confluent model $M' = (W, \leq, (R; \leq), V)$ that satisfies $M, w \models \psi$ iff $M', w \models \psi$ for all $w \in W$ and formulas ψ.

(Algebraic) semantics of IL_μ: monotonicity and confluence properties

A key property of intuitionistic Kripke semantics is monotonicity: if $v \geq w$ and $w \models \varphi$, then $v \models \varphi$.

To obtain monotonicity for IL_μ, models need to satisfy triangle confluence: if $w \leq vRu$ then wRu.

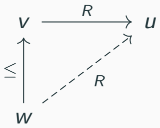

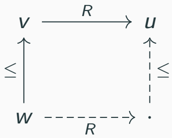

For $\Box$-formulas, forth-down confluence suffices: if $w \leq vRu$ then $wRs \leq u$ for some s.

Lemma

A $\Box$-formula φ is valid on all forth-down confluent models iff it is valid on all triangle confluent models.

Proof. Any forth-down confluent model $M = (W, \leq, R, V)$ induces a triangle confluent model $M' = (W, \leq, (R; \leq), V)$ that satisfies $M, w \models \psi$ iff $M', w \models \psi$ for all $w \in W$ and formulas ψ.

As $\dashv$-formulas are not monotone for the weaker condition, we obtain that $\dashv$ indeed cannot be expressed in terms of $\Box$.

Game semantics for IL_μ

Game semantics for IL_μ: the evaluation game

Given a model $M = (W, \leq, R, V)$ and a clean formula φ, we define an evaluation game $\mathcal{E}(\varphi, M)$ between $\exists$ and $\forall$ to determine which states $s \in W$ satisfy $M, s \models \varphi$.

Intuition: at position $(\psi, s, +)$, $\exists$ wants to show that s satisfies ψ, while $\forall$ wants to show the converse.

Game semantics for IL_μ: the evaluation game

Given a model $M = (W, \leq, R, V)$ and a clean formula φ, we define an evaluation game $\mathcal{E}(\varphi, M)$ between $\exists$ and $\forall$ to determine which states $s \in W$ satisfy $M, s \models \varphi$.

Intuition: at position $(\psi, s, +)$, $\exists$ wants to show that s satisfies ψ, while $\forall$ wants to show the converse.

Position	Player	Admissible moves
$(\top, s, +)$	$\forall$	$\emptyset$
$(\bot, s, +)$	$\exists$	$\emptyset$
$(P, s, +)$, $P \notin BV(\psi)$, $s \in V(P)$	$\forall$	$\emptyset$
$(P, s, +)$, $P \notin BV(\psi)$, $s \notin V(P)$	$\exists$	$\emptyset$
$(\varphi_1 \wedge \varphi_2, s, +)$	$\forall$	$\{(\varphi_i, s, +) : i = 1, 2\}$
$(\varphi_1 \vee \varphi_2, s, +)$	$\exists$	$\{(\varphi_i, s, +) : i = 1, 2\}$
$(\varphi_1 \to \varphi_2, s, +)$	$\forall$	$\{(\varphi_1 \to \varphi_2, s, t, +) : s \leq t\}$
$(\varphi_1 \to \varphi_2, s, t, +)$	$\exists$	$\{(\varphi_1, t, -), (\varphi_2, t, +)\}$
$(\varphi_1 \prec\!\!\!\!\!\!\rhd \varphi_2, s, +)$	$\forall$	$\{(\varphi_1 \prec\!\!\!\!\!\!\rhd \varphi_2, s, t, +) : sRt\}$
$(\varphi_1 \prec\!\!\!\!\!\!\rhd \varphi_2, s, t, +)$	$\exists$	$\{(\varphi_1, t, -), (\varphi_2, t, +)\}$
$(\sigma_X X.\delta_X, s, +)$	-	$\{(\delta_X, s, +)\}$
$(X, s, +)$, $X \in BV(\psi)$	-	$\{(\delta_X, s, +)\}$

Game semantics for IL_μ: the evaluation game

Given a model $M = (W, \leq, R, V)$ and a clean formula φ, we define an evaluation game $\mathcal{E}(\varphi, M)$ between $\exists$ and $\forall$ to determine which states $s \in W$ satisfy $M, s \models \varphi$.

Intuition: at position $(\psi, s, +)$, $\exists$ wants to show that s satisfies ψ, while $\forall$ wants to show the converse.

Position	Player	Admissible moves
$(\top, s, +)$	$\forall$	$\emptyset$
$(\bot, s, +)$	$\exists$	$\emptyset$
$(P, s, +)$, $P \notin BV(\psi)$, $s \in V(P)$	$\forall$	$\emptyset$
$(P, s, +)$, $P \notin BV(\psi)$, $s \notin V(P)$	$\exists$	$\emptyset$
$(\varphi_1 \wedge \varphi_2, s, +)$	$\forall$	$\{(\varphi_i, s, +) : i = 1, 2\}$
$(\varphi_1 \vee \varphi_2, s, +)$	$\exists$	$\{(\varphi_i, s, +) : i = 1, 2\}$
$(\varphi_1 \to \varphi_2, s, +)$	$\forall$	$\{(\varphi_1 \to \varphi_2, s, t, +) : s \leq t\}$
$(\varphi_1 \to \varphi_2, s, t, +)$	$\exists$	$\{(\varphi_1, t, -), (\varphi_2, t, +)\}$
$(\varphi_1 \dashv \varphi_2, s, +)$	$\forall$	$\{(\varphi_1 \dashv \varphi_2, s, t, +) : sRt\}$
$(\varphi_1 \dashv \varphi_2, s, t, +)$	$\exists$	$\{(\varphi_1, t, -), (\varphi_2, t, +)\}$
$(\sigma_X X.\delta_X, s, +)$	-	$\{(\delta_X, s, +)\}$
$(X, s, +)$, $X \in BV(\psi)$	-	$\{(\delta_X, s, +)\}$

For negative positions $(\varphi, s, -)$, we swap the roles of $\exists$ and $\forall$. We call $+$ or $-$ the parity of a position.

Game semantics for IL_μ: winning conditions

Let q be a position. A play of $\mathcal{E}(\varphi, M)@q$ is a sequence ρ of positions following the rules of $\mathcal{E}(\varphi, M)$ such that

1. $\rho_0 = q$,
2. ρ is either infinite or ending in a position with no admissible moves.

Game semantics for IL_μ: winning conditions

Let q be a position. A play of $\mathcal{E}(\varphi, M)@q$ is a sequence ρ of positions following the rules of $\mathcal{E}(\varphi, M)$ such that

1. $\rho_0 = q$,
2. ρ is either infinite or ending in a position with no admissible moves.

Who wins a play ρ? We want winning conditions such that: $M, s \models \varphi$ iff $\exists$ has a winning strategy in $\mathcal{E}(\varphi, M)@(\varphi, s, +)$.

Game semantics for IL_μ: winning conditions

Let q be a position. A play of $\mathcal{E}(\varphi, M)@q$ is a sequence ρ of positions following the rules of $\mathcal{E}(\varphi, M)$ such that

1. $\rho_0 = q$,
2. ρ is either infinite or ending in a position with no admissible moves.

Who wins a play ρ? We want winning conditions such that: $M, s \models \varphi$ iff $\exists$ has a winning strategy in $\mathcal{E}(\varphi, M)@(\varphi, s, +)$.

Lemma

Let ρ an infinite play of $\mathcal{E}(\varphi, M)@(\varphi, s, +)$. Then there is a unique, outermost $X_\rho \in BV(\varphi)$ that occurs infinitely often in ρ. Moreover, there is a parity $\bullet_\rho \in \{+, -\}$ such that every position in ρ with formula X has parity $\bullet_\rho$.

Let q be a position. A play of $\mathcal{E}(\varphi, M)@q$ is a sequence ρ of positions following the rules of $\mathcal{E}(\varphi, M)$ such that

1. $\rho_0 = q$,
2. ρ is either infinite or ending in a position with no admissible moves.

Who wins a play ρ? We want winning conditions such that: $M, s \models \varphi$ iff $\exists$ has a winning strategy in $\mathcal{E}(\varphi, M)@(\varphi, s, +)$.

Lemma

Let ρ an infinite play of $\mathcal{E}(\varphi, M)@(\varphi, s, +)$. Then there is a unique, outermost $X_\rho \in BV(\varphi)$ that occurs infinitely often in ρ. Moreover, there is a parity $\bullet_\rho \in \{+, -\}$ such that every position in ρ with formula X has parity $\bullet_\rho$.

Winning conditions: Finite plays are lost by the player who got stuck. An infinite play ρ is won by $\exists$ if $(\sigma_{X_\rho}, \bullet_\rho) \in \{(\nu, +), (\mu, -)\}$, and won by $\forall$ otherwise.

Game semantics for IL_μ: determinacy and adequacy

The game $\mathcal{E}(\varphi, M)$ is a parity game and therefore positionally determined: for every position q, either $\exists$ or $\forall$ has a winning strategy in $\mathcal{E}(\varphi, M)@q$ that only depends on the current position.

Game semantics for IL_μ: determinacy and adequacy

The game $\mathcal{E}(\varphi, M)$ is a parity game and therefore positionally determined: for every position q, either $\exists$ or $\forall$ has a winning strategy in $\mathcal{E}(\varphi, M)@q$ that only depends on the current position.

Theorem

For any clean formula φ and pointed model (M, s), we have $M, s \models \varphi$ iff $\exists$ has a (positional) winning strategy in $\mathcal{E}(\varphi, M)@(\varphi, s, +)$.

Game semantics for IL_μ: determinacy and adequacy

The game $\mathcal{E}(\varphi, M)$ is a parity game and therefore positionally determined: for every position q, either $\exists$ or $\forall$ has a winning strategy in $\mathcal{E}(\varphi, M)@q$ that only depends on the current position.

Theorem

For any clean formula φ and pointed model (M, s), we have $M, s \models \varphi$ iff $\exists$ has a (positional) winning strategy in $\mathcal{E}(\varphi, M)@(\varphi, s, +)$.

Proof. By induction on φ. We adapt the proof for the classical case, making use of two observations:

1. For any position $(\psi, t, +)$: $(\psi, t, +)$ is winning for $\exists$ iff $(\psi, t, -)$ is winning for $\forall$.
2. Every bound variable has a fixed parity in $\mathcal{E}(\varphi, M)@(\varphi, s, +)$.

Game semantics for IL_μ: determinacy and adequacy

The game $\mathcal{E}(\varphi, M)$ is a parity game and therefore positionally determined: for every position q, either $\exists$ or $\forall$ has a winning strategy in $\mathcal{E}(\varphi, M)@q$ that only depends on the current position.

Theorem

For any clean formula φ and pointed model (M, s), we have $M, s \models \varphi$ iff $\exists$ has a (positional) winning strategy in $\mathcal{E}(\varphi, M)@(\varphi, s, +)$.

Proof. By induction on φ. We adapt the proof for the classical case, making use of two observations:

1. For any position $(\psi, t, +)$: $(\psi, t, +)$ is winning for $\exists$ iff $(\psi, t, -)$ is winning for $\forall$.
2. Every bound variable has a fixed parity in $\mathcal{E}(\varphi, M)@(\varphi, s, +)$.

Remark. There is a preprint by Pacheco (2023) where similar (independently developed) game semantics for an intuitionistic version of the modal μ-calculus are used to show a collapse to modal logic over intuitionistic S5 frames.

Guardedness

Guardedness: definition

In the classical μ-calculus, it is well known that every formula is equivalent to a guarded one.

Given a formula φ and variable X, we call X **guarded in φ** if every occurrence of X in φ is in the scope of some −ɜ-operator. A formula φ is **guarded** if for every subformula $\sigma X.\psi$ of φ, X is guarded in ψ.

Guardedness: definition

In the classical μ-calculus, it is well known that every formula is equivalent to a guarded one.

Given a formula φ and variable X, we call X guarded in φ if every occurrence of X in φ is in the scope of some $\prec\!\!\!\!\exists$-operator. A formula φ is guarded if for every subformula $\sigma X.\psi$ of φ, X is guarded in ψ.

We use two results to show that every formula of IL_μ is equivalent to a guarded one:

Theorem (Ruitenburg, 1984)

Let φ be a formula of IPC and X a propositional letter such that X is positive in φ. Define $\varphi^0_X := X$ and $\varphi^{n+1}_X := \varphi[\varphi^n_X/X]$. Then there exists an N such that $\varphi^N_X \equiv \varphi^{N+1}_X$.

In particular, we have $\mu X.\varphi \equiv \varphi^N_X[\bot/X]$ and $\nu X.\varphi \equiv \varphi^N_X[\top/X]$.

Guardedness: definition

In the classical μ-calculus, it is well known that every formula is equivalent to a guarded one.

Given a formula φ and variable X, we call X guarded in φ if every occurrence of X in φ is in the scope of some $\prec\!\!\exists$-operator. A formula φ is guarded if for every subformula $\sigma X.\psi$ of φ, X is guarded in ψ.

We use two results to show that every formula of IL_μ is equivalent to a guarded one:

Theorem (Ruitenburg, 1984)

Let φ be a formula of IPC and X a propositional letter such that X is positive in φ. Define $\varphi_X^0 := X$ and $\varphi_X^{n+1} := \varphi[\varphi_X^n/X]$. Then there exists an N such that $\varphi_X^N \equiv \varphi_X^{N+1}$.

In particular, we have $\mu X.\varphi \equiv \varphi_X^N[\bot/X]$ and $\nu X.\varphi \equiv \varphi_X^N[\top/X]$.

Lemma ("the golden lemma of the μ-calculus")

For any $\varphi(X, Y)$ with X and Y positive, we have

$$\sigma X.\sigma Y.\varphi(X, Y) \equiv \sigma X.\varphi(X, X) \equiv \sigma Y.\sigma X.\varphi(X, Y).$$

Proof. Straightforward by the game semantics!

Guardedness: proof

Theorem

Every IL_μ*-formula* φ *is equivalent to a guarded one.*

Guardedness: proof

Theorem

Every IL_μ-formula φ is equivalent to a guarded one.

Proof. By induction on φ. Suppose φ is of the form $\mu X.\psi$ with ψ guarded. Due to unfolding, we may assume that all unguarded occurrences of X in ψ are outside scope of fixpoint operators.

Guardedness: proof

Theorem

Every IL_μ*-formula* φ *is equivalent to a guarded one.*

Proof. By induction on φ. Suppose φ is of the form $\mu X.\psi$ with ψ guarded. Due to unfolding, we may assume that all unguarded occurrences of X in ψ are outside scope of fixpoint operators.

Let ζ be obtained from ψ by replacing every unguarded occurrence of X by X_0 and every guarded occurrence of X by X_1. Then $\mu X.\psi \equiv \mu X_0.\mu X_1.\zeta \equiv \mu X_1.\mu X_0.\zeta$.

Guardedness: proof

Theorem

Every IL_μ*-formula* φ *is equivalent to a guarded one.*

Proof. By induction on φ. Suppose φ is of the form $\mu X.\psi$ with ψ guarded. Due to unfolding, we may assume that all unguarded occurrences of X in ψ are outside scope of fixpoint operators.

Let ζ be obtained from ψ by replacing every unguarded occurrence of X by X_0 and every guarded occurrence of X by X_1. Then $\mu X.\psi \equiv \mu X_0.\mu X_1.\zeta \equiv \mu X_1.\mu X_0.\zeta$.

Let $\hat{\zeta}$ be obtained from ζ by replacing each (maximal) fixpoint or modal subformula χ of ψ by a fresh propositional letter P_χ; by construction, no such χ contains the variable X_0.

Guardedness: proof

Theorem

Every IL_μ*-formula* φ *is equivalent to a guarded one.*

Proof. By induction on φ. Suppose φ is of the form $\mu X.\psi$ with ψ guarded. Due to unfolding, we may assume that all unguarded occurrences of X in ψ are outside scope of fixpoint operators.

Let ζ be obtained from ψ by replacing every unguarded occurrence of X by X_0 and every guarded occurrence of X by X_1. Then $\mu X.\psi \equiv \mu X_0.\mu X_1.\zeta \equiv \mu X_1.\mu X_0.\zeta$.

Let $\hat{\zeta}$ be obtained from ζ by replacing each (maximal) fixpoint or modal subformula χ of ψ by a fresh propositional letter P_χ; by construction, no such χ contains the variable X_0.

Then $\hat{\zeta}$ is an IPC formula, so by Ruitenburg's theorem there exists an N such that $\hat{\zeta}^N_{X_0} \equiv \hat{\zeta}^{N+1}_{X_0}$.

Theorem

Every IL_μ-formula φ is equivalent to a guarded one.

Proof. By induction on φ. Suppose φ is of the form $\mu X.\psi$ with ψ guarded. Due to unfolding, we may assume that all unguarded occurrences of X in ψ are outside scope of fixpoint operators.

Let ζ be obtained from ψ by replacing every unguarded occurrence of X by X_0 and every guarded occurrence of X by X_1. Then $\mu X.\psi \equiv \mu X_0.\mu X_1.\zeta \equiv \mu X_1.\mu X_0.\zeta$.

Let $\hat{\zeta}$ be obtained from ζ by replacing each (maximal) fixpoint or modal subformula χ of ψ by a fresh propositional letter P_χ; by construction, no such χ contains the variable X_0.

Then $\hat{\zeta}$ is an IPC formula, so by Ruitenburg's theorem there exists an N such that $\hat{\zeta}^N_{X_0} \equiv \hat{\zeta}^{N+1}_{X_0}$.

Let $\overline{\zeta} := \zeta^N_{X_0}[\bot/X_0]$. As none of the χ contains X_0, note that $\zeta^N_{X_0}$ is identical to the formula obtained from $\hat{\zeta}^N_{X_0}$ by substituting for each P_χ its corresponding subformula χ. Since $\hat{\zeta}^N_{X_0} \equiv \hat{\zeta}^{N+1}_{X_0}$, it follows that $\zeta^N_{X_0} \equiv \zeta^{N+1}_{X_0}$, which implies $\mu X_0.\zeta \equiv \zeta^N_{X_0}[\bot/X_0]$.

Guardedness: proof

Theorem

Every IL_μ-formula φ is equivalent to a guarded one.

Proof. By induction on φ. Suppose φ is of the form $\mu X.\psi$ with ψ guarded. Due to unfolding, we may assume that all unguarded occurrences of X in ψ are outside scope of fixpoint operators.

Let ζ be obtained from ψ by replacing every unguarded occurrence of X by X_0 and every guarded occurrence of X by X_1. Then $\mu X.\psi \equiv \mu X_0.\mu X_1.\zeta \equiv \mu X_1.\mu X_0.\zeta$.

Let $\hat{\zeta}$ be obtained from ζ by replacing each (maximal) fixpoint or modal subformula χ of ψ by a fresh propositional letter P_χ; by construction, no such χ contains the variable X_0.

Then $\hat{\zeta}$ is an IPC formula, so by Ruitenburg's theorem there exists an N such that $\hat{\zeta}^N_{X_0} \equiv \hat{\zeta}^{N+1}_{X_0}$.

Let $\overline{\zeta} := \zeta^N_{X_0}[\bot/X_0]$. As none of the χ contains X_0, note that $\zeta^N_{X_0}$ is identical to the formula obtained from $\hat{\zeta}^N_{X_0}$ by substituting for each P_χ its corresponding subformula χ. Since $\hat{\zeta}^N_{X_0} \equiv \hat{\zeta}^{N+1}_{X_0}$, it follows that $\zeta^N_{X_0} \equiv \zeta^{N+1}_{X_0}$, which implies $\mu X_0.\zeta \equiv \zeta^N_{X_0}[\bot/X_0]$.

Now $\zeta^N_{X_0}[\bot/X_0]$ only contains guarded occurrences of X_1, and $\mu X_1.\zeta^N_{X_0}[\bot/X_0] \equiv \mu X_1.\mu X_0.\zeta \equiv \mu X.\psi$.

Guardedness: proof

Theorem

Every IL_μ*-formula* φ *is equivalent to a guarded one.*

Proof. By induction on φ. Suppose φ is of the form $\mu X.\psi$ with ψ guarded. Due to unfolding, we may assume that all unguarded occurrences of X in ψ are outside scope of fixpoint operators.

Let ζ be obtained from ψ by replacing every unguarded occurrence of X by X_0 and every guarded occurrence of X by X_1. Then $\mu X.\psi \equiv \mu X_0.\mu X_1.\zeta \equiv \mu X_1.\mu X_0.\zeta$.

Let $\hat{\zeta}$ be obtained from ζ by replacing each (maximal) fixpoint or modal subformula χ of ψ by a fresh propositional letter P_χ; by construction, no such χ contains the variable X_0.

Then $\hat{\zeta}$ is an IPC formula, so by Ruitenburg's theorem there exists an N such that $\hat{\zeta}^N_{X_0} \equiv \hat{\zeta}^{N+1}_{X_0}$.

Let $\overline{\zeta} := \zeta^N_{X_0}[\bot/X_0]$. As none of the χ contains X_0, note that $\zeta^N_{X_0}$ is identical to the formula obtained from $\hat{\zeta}^N_{X_0}$ by substituting for each P_χ its corresponding subformula χ. Since $\hat{\zeta}^N_{X_0} \equiv \hat{\zeta}^{N+1}_{X_0}$, it follows that $\zeta^N_{X_0} \equiv \zeta^{N+1}_{X_0}$, which implies $\mu X_0.\zeta \equiv \zeta^N_{X_0}[\bot/X_0]$.

Now $\zeta^N_{X_0}[\bot/X_0]$ only contains guarded occurrences of X_1, and $\mu X_1.\zeta^N_{X_0}[\bot/X_0] \equiv \mu X_1.\mu X_0.\zeta \equiv \mu X.\psi$.

The ν-case is completely analoguous. □

A non-wellfounded proof system for IL_μ

A non-wellfounded proof system for IL_μ: rules

We consider sequents $\Gamma \Rightarrow \Delta$ of finite sets of clean formulas with the standard interpretation:

$$\bigwedge \Gamma \to \bigvee \Delta.$$

We define a non-wellfounded sequent calculus nwIL_μ as follows.

A non-wellfounded proof system for IL_μ: rules

We consider sequents $\Gamma \Rightarrow \Delta$ of finite sets of clean formulas with the standard interpretation:

$$\bigwedge \Gamma \to \bigvee \Delta.$$

We define a non-wellfounded sequent calculus nwIL_μ as follows.

For the propositional rules, we use standard multi-conclusion rules for IPC:

$$\frac{}{\Gamma, A \Rightarrow A, \Delta}\ \mathrm{id} \qquad \frac{}{\Gamma, \bot \Rightarrow \Delta}\ \bot$$

$$\frac{\Gamma, A, B \Rightarrow \Delta}{\Gamma, A \wedge B \Rightarrow \Delta}\ \wedge\mathrm{L} \qquad \frac{\Gamma \Rightarrow A, \Delta \quad \Gamma \Rightarrow B, \Delta}{\Gamma \Rightarrow A \wedge B, \Delta}\ \wedge\mathrm{R}$$

$$\frac{\Gamma, A \Rightarrow \Delta \quad \Gamma, B \Rightarrow \Delta}{\Gamma, A \vee B \Rightarrow \Delta}\ \vee\mathrm{L} \qquad \frac{\Gamma \Rightarrow A, B, \Delta}{\Gamma \Rightarrow A \vee B, \Delta}\ \vee\mathrm{R}$$

$$\frac{\Gamma, A \to B \Rightarrow A, \Delta \quad \Gamma, B \Rightarrow \Delta}{\Gamma, A \to B \Rightarrow \Delta}\ {\to}\mathrm{L} \qquad \frac{\Gamma, A \Rightarrow B}{\Gamma \Rightarrow A \to B, \Delta}\ {\to}\mathrm{R}$$

A non-wellfounded proof system for IL_μ: rules

Consider the following sound rule for the modality $\rhd$:

$$\frac{A \Rightarrow B, C \quad D, A \Rightarrow B}{\Gamma, C \rhd D \Rightarrow A \rhd B, \Delta}\ \rhd_1$$

A non-wellfounded proof system for IL_μ: rules

Consider the following sound rule for the modality $\rhd$:

$$\frac{A \Rightarrow B, C \quad D, A \Rightarrow B}{\Gamma, C \rhd D \Rightarrow A \rhd B, \Delta} \rhd_1$$

For completeness, we generalize it to the following:

$$\frac{\{\mathcal{D}_j, A \Rightarrow B, \mathcal{C}_j\}_{j \leq 2^n}}{\Gamma, \{C_i \rhd D_i\}_{i \leq n} \Rightarrow A \rhd B, \Delta} \rhd_n$$

where $n \geq 0$, and the sets $\mathcal{D}_1, \dots, \mathcal{D}_{2^n}$ and $\mathcal{C}_1, \dots, \mathcal{C}_{2^n}$ enumerate the subsets of $\{D_1, \dots, D_n\}$ and $\{C_1, \dots, C_n\}$, respectively, such that

$$D_i \in \mathcal{D}_j \text{ if and only if } C_i \notin \mathcal{C}_j.$$

A non-wellfounded proof system for IL_μ: rules

Consider the following sound rule for the modality $\rhd$:

$$\frac{A \Rightarrow B, C \quad D, A \Rightarrow B}{\Gamma, C \rhd D \Rightarrow A \rhd B, \Delta} \ \rhd_1$$

For completeness, we generalize it to the following:

$$\frac{\{\mathcal{D}_j, A \Rightarrow B, \mathcal{C}_j\}_{j \leq 2^n}}{\Gamma, \{C_i \rhd D_i\}_{i \leq n} \Rightarrow A \rhd B, \Delta} \ \rhd_n$$

where $n \geq 0$, and the sets $\mathcal{D}_1, \ldots, \mathcal{D}_{2^n}$ and $\mathcal{C}_1, \ldots, \mathcal{C}_{2^n}$ enumerate the subsets of $\{D_1, \ldots, D_n\}$ and $\{C_1, \ldots, C_n\}$, respectively, such that

$$D_i \in \mathcal{D}_j \text{ if and only if } C_i \notin \mathcal{C}_j.$$

Remark. A (sound) single-conclusion version

$$\frac{A \Rightarrow C_1 \quad D_1, A \Rightarrow C_2 \quad \ldots \quad D_1, \ldots, D_n, A \Rightarrow B}{\Gamma, \{C_i \rhd D_i\}_{i \leq n} \Rightarrow A \rhd B}$$

is not complete, as it cannot derive the valid formula $\varphi \rhd \chi \to (\psi \rhd \chi \to (\varphi \vee \psi) \rhd \chi)$.

For the fixpoint rules, we assume a derivation comes with a mapping $X \mapsto (\sigma_X, \delta_X)$ that associates every bound variable with its corresponding binder and fixpoint formula.

For $\sigma \in \{\mu, \nu\}$, we have the rules

$$\frac{\Gamma, \delta \Rightarrow \Delta}{\Gamma, \sigma X.\delta \Rightarrow \Delta}\ \sigma\mathrm{L} \quad \frac{\Gamma \Rightarrow \delta, \Delta}{\Gamma \Rightarrow \sigma X.\delta, \Delta}\ \sigma\mathrm{R}$$

$$\frac{\Gamma, \delta_X \Rightarrow \Delta}{\Gamma, X \Rightarrow \Delta}\ X\mathrm{L} \qquad \frac{\Gamma \Rightarrow \delta_X, \Delta}{\Gamma \Rightarrow X, \Delta}\ X\mathrm{R}$$

This concludes the rules of nwIL_μ.

A non-wellfounded proof system for IL_μ: derivations and proofs

A derivation T in nwIL_μ is a finite or infinite tree labelled according to the rules of nwIL_μ.

A non-wellfounded proof system for IL_μ: derivations and proofs

A derivation T in nwIL_μ is a finite or infinite tree labelled according to the rules of nwIL_μ.

Given a path ρ through T, a trace on ρ is a sequence $(\varphi_i, \bullet_i)$ of formulas and their parity (left/right) following the principal-residual relation.

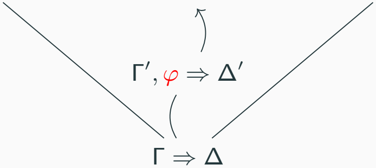

A non-wellfounded proof system for IL_μ: derivations and proofs

A derivation T in nwIL$_\mu$ is a finite or infinite tree labelled according to the rules of nwIL$_\mu$.

Given a path ρ through T, a trace on ρ is a sequence $(\varphi_i, \bullet_i)$ of formulas and their parity (left/right) following the principal-residual relation.

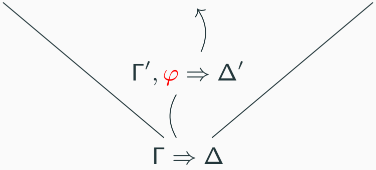

Similarly to plays, each (non-stagnating) trace τ will contain a unique outermost bound variable X that occurs infinitely often. This X is bound by either μ or ν, and its parity is either left or right.

A non-wellfounded proof system for IL_μ: derivations and proofs

A derivation T in nwIL$_\mu$ is a finite or infinite tree labelled according to the rules of nwIL$_\mu$.

Given a path ρ through T, a trace on ρ is a sequence $(\varphi_i, \bullet_i)$ of formulas and their parity (left/right) following the principal-residual relation.

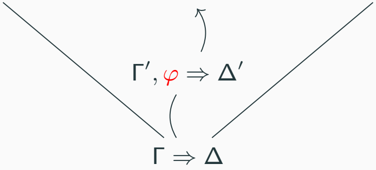

Similarly to plays, each (non-stagnating) trace τ will contain a unique outermost bound variable X that occurs infinitely often. This X is bound by either μ or ν, and its parity is either left or right.

A derivation T is a proof in nwIL$_\mu$ if

1. every leaf of T is labelled by an axiom;
2. every infinite path of T has either a left μ-trace or a right ν-trace.

A non-wellfounded proof system for IL_μ: soundness and completeness

Theorem

If φ is provable in nwIL_μ *then it is valid on triangle models.*

Theorem

Every guarded formula valid on triangle models is provable in nwIL_μ*.*

Proof. By a game-theoretic argument similar to that of Niwiński and Walukiewicz (1996).

A non-wellfounded proof system for IL_μ: soundness and completeness

Theorem

If φ is provable in nwIL_μ *then it is valid on triangle models.*

Theorem

Every guarded formula valid on triangle models is provable in nwIL_μ.

Proof. By a game-theoretic argument similar to that of Niwiński and Walukiewicz (1996).

- Given a sequent σ, construct a determined two-player game between Prover and Refuter.

A non-wellfounded proof system for IL_μ: soundness and completeness

Theorem

If φ is provable in nwIL_μ *then it is valid on triangle models.*

Theorem

Every guarded formula valid on triangle models is provable in nwIL_μ.

Proof. By a game-theoretic argument similar to that of Niwiński and Walukiewicz (1996).

- Given a sequent σ, construct a determined two-player game between Prover and Refuter.
- The game is played on a proof search tree. In this tree, the non-invertible rules $\to\mathrm{R}$ and $\dashv_n$ may only be applied once the invertible rules have been applied to a sufficient degree.

A non-wellfounded proof system for IL_μ: soundness and completeness

Theorem

If φ is provable in nwIL_μ then it is valid on triangle models.

Theorem

Every guarded formula valid on triangle models is provable in nwIL_μ.

Proof. By a game-theoretic argument similar to that of Niwiński and Walukiewicz (1996).

- Given a sequent σ, construct a determined two-player game between Prover and Refuter.
- The game is played on a proof search tree. In this tree, the non-invertible rules $\to\mathrm{R}$ and $\triangleleft_n$ may only be applied once the invertible rules have been applied to a sufficient degree.
- A winning strategy for Prover corresponds to a proof of σ.

A non-wellfounded proof system for IL_μ: soundness and completeness

Theorem

If φ is provable in nwIL_μ *then it is valid on triangle models.*

Theorem

Every guarded formula valid on triangle models is provable in nwIL_μ*.*

Proof. By a game-theoretic argument similar to that of Niwiński and Walukiewicz (1996).

- Given a sequent σ, construct a determined two-player game between Prover and Refuter.
- The game is played on a proof search tree. In this tree, the non-invertible rules $\to\mathrm{R}$ and $\dashv_n$ may only be applied once the invertible rules have been applied to a sufficient degree.
- A winning strategy for Prover corresponds to a proof of σ.
- From a winning strategy for Refuter, we construct a (pre-)countermodel M for σ by treating a premise of $\dashv_n$ as modal successor and a premise of $\to$R as intuitionistic successor.

A non-wellfounded proof system for IL_μ: soundness and completeness

Theorem

If φ is provable in nwIL_μ *then it is valid on triangle models.*

Theorem

Every guarded formula valid on triangle models is provable in nwIL_μ*.*

Proof. By a game-theoretic argument similar to that of Niwiński and Walukiewicz (1996).

- Given a sequent σ, construct a determined two-player game between Prover and Refuter.
- The game is played on a proof search tree. In this tree, the non-invertible rules $\to\mathrm{R}$ and $\dashv_n$ may only be applied once the invertible rules have been applied to a sufficient degree.
- A winning strategy for Prover corresponds to a proof of σ.
- From a winning strategy for Refuter, we construct a (pre-)countermodel M for σ by treating a premise of $\dashv_n$ as modal successor and a premise of $\to$R as intuitionistic successor.
- We make M satisfy triangle confluence by replacing the modal relation R by the composition $\leq; R$.

Future work

Future work

We have introduced the logic IL_μ and provided adequate game semantics and a complete, analytic non-wellfounded proof system for guarded formulas.

Future work

We have introduced the logic IL_μ and provided adequate game semantics and a complete, analytic non-wellfounded proof system for guarded formulas.

The next steps:

- Prove completeness for unguarded formulas.
- Use the analytic, non-wellfounded system to obtain a cyclic system with annotations.
- Prove the bounded model property, and thereby decidability.

Future work

We have introduced the logic IL_μ and provided adequate game semantics and a complete, analytic non-wellfounded proof system for guarded formulas.

The next steps:

- Prove completeness for unguarded formulas.
- Use the analytic, non-wellfounded system to obtain a cyclic system with annotations.
- Prove the bounded model property, and thereby decidability.

Bigger steps:

- Study proof-theoretic properties such as cut-elimination, interpolation...
- Consider different frame conditions.
- Adding diamonds??