Logic and Computation II

Part 5. Modal μ-calculus

Kazuyuki Tanaka

BIMSA

April 3, 2025

北京雁栖湖
应用数学研究院
Beijing Institute
of Mathematical
Sciences and Applications

Logic and Computation II

- **Part 4. Modal logic**
- **Part 5. Modal μ-calculus**
- **Part 6. Automata on infinite objects**
- **Part 7. Recursion-theoretic hierarchies**

Part 5. Schedule (tentative)

- March 27, (1) Introduction to modal μ-calculus and monadic second-order logic
- April 1, (2) Basics of modal μ-calculus
- April 3, (3) The adequacy theorem

Recap: basics of modal mu-Calculus

The formulas of the modal μ-calculus, $\mathcal{L}_\mu$, are defined as follows:

$$\varphi ::= p \mid \neg p \mid X \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \Box\varphi \mid \Diamond\varphi \mid \mu X.\varphi \mid \nu X.\varphi,$$

where p is an atomic proposition and X is a (proposional) variable. For convenience, the negation of a formula is introduced by the following rules:

$$\begin{array}{rclcrcl} \neg\neg\varphi & \equiv & \varphi, & & \neg\Box\varphi & \equiv & \Diamond\neg\varphi, \\ \neg(\varphi\wedge\psi) & \equiv & \neg\varphi\vee\neg\psi, & & \neg(\varphi\vee\psi) & \equiv & \neg\varphi\wedge\neg\psi, \\ \neg\mu X.\varphi & \equiv & \nu X.\neg\varphi[\neg X/X], & & & & \end{array}$$

where $\varphi[\neg X/X]$ is obtained by replacing every free occurrence of X in φ with $\neg X$.

The truth valuation $V(\varphi) = \{s : M, s \models \varphi\}$, also written as $||\varphi||^M$, is defined as in ordinary modal logic except for $\eta X.\theta(X)$. Regarding the variable X as an atomic proposition with $V(X) = S \subseteq W$, we define a monotone function $\Psi(S)$ as follows:

$$\Psi(S) := ||\theta(X)||^M_{V(X)=S}.$$

Then, the valuation of $\eta X.\theta(X)$ are given as follows:

$$||\mu X.\theta(X)||^M := \bigcap\{S \subseteq W : \Psi(S) \subseteq S\},\ ||\nu X.\theta(X)||^M := \bigcup\{S \subseteq W : \Psi(S) \supseteq S\}$$

The Evaluation Game

A Kripke model M is fixed. For a state s and a formula φ, we consider the **evaluation game** $\mathcal{E}(M, s, \varphi)$. In this game, two players $\exists$ (female) and $\forall$ (male) decompose the formula while moving a token from s over the frame (much like a board game).

We first consider multi-modal logic with a model $M = (W, (R_i)_{i \in I}, v)$. A **position** in the game $\mathcal{E}(M, s_0, \varphi)$ is a pair (s, ψ) where $s \in W$ and ψ is a subformula of φ. The rules are as follows: (s_0, φ) is the initial position.

- At position $(s, \psi \vee \psi')$, it is $\exists$'s turn to choose either (s, ψ) or (s, ψ').
- At position $(s, \psi \wedge \psi')$, $\forall$ chooses a next position (s, ψ) or (s, ψ').
- At position $(s, \Box_i \psi)$, $\forall$ chooses a next position (t, ψ) such that there is an i-labeled edge from s to t. (If no such t exists, $\forall$ loses.)
- At $(s, \Diamond_i \psi)$, $\exists$ chooses such a position (t, ψ). (If no such t exists, $\exists$ loses.)

When a terminal position (s, p) or $(s, \neg p)$ is reached, $\exists$ **wins** if $s \in v(p)$ or $s \notin v(p)$, resp.

Adequacy Theorem

$\exists$ has a winning strategy in $\mathcal{E}(M, s, \varphi) \iff M, s \models \varphi$

Extending the game to modal μ-calculus

For simplicity, we assume that any two fixpoint variables appearing in a formula are distinct. We add the following rules for new positions:

- A position $(s, \eta X.\theta)$ automatically changes to (s, θ).
- Later, when the play reaches (t, X), the game jumps back to $(t, \eta X.\theta)$.

This means that the game does not need to terminate in a finite way. If the play terminates in a finite number of moves, the winning condition is the same as in the modal logic case.

If the play is infinite, then the winner is decided by the form of the *outermost* (largest) fixed point formula appearing infinitely often:

- If the outmost fixpoint is of the form $\mu X.\varphi$, then $\forall$ wins.
- If it is of the form $\nu X.\varphi$, then $\exists$ wins.

Example

The following are equivalent:

1. $M, s \models \mu X.p \vee \Diamond_i X$.
2. There exists a path from s (via R_i edges) reaching a state where p holds.
3. $\exists$ has a winning strategy in $\mathcal{E}(M, s, \mu X.p \vee \Diamond_i X)$.

Problem 1.

Assume a frame M with edges labeled a and b.

(1) Write a modal μ-calculus formula that expresses:
"From state s, there is a path that contains infinitely many a-labeled edges."

(2) Write a modal μ-calculus formula that expresses:
"From state s, there is a path that contains infinitely many a-labeled edges, and there is no path containing infinitely many b-labeled edges."

Answers:

(1) $\nu X.\mu Y.\Diamond_a X \vee \Diamond_b Y$.

(2) $(\nu X.\mu Y.\Diamond_a X \vee \Diamond_b Y) \wedge (\mu Y.\nu X.\Box_a X \wedge \Box_b Y)$.

Example

Consider $\mathcal{M}$ as follows, where the only atomic proposition is p, and $v(p) = W$ (i.e., p is always true). Analyze the evaluation game.

$$\mathcal{E}(\mathcal{M}, 0, \mu Y.\nu Z.\Box_a((\Diamond_b \top \vee Y) \wedge Z)).$$

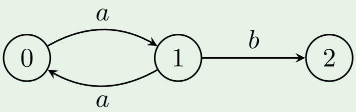

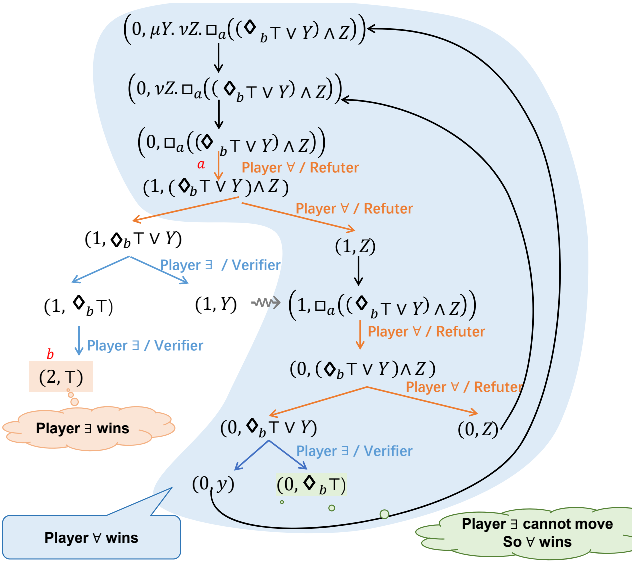

§5.3. Fundamental theorems in $\mathcal{L}_\mu$

Adequacy Theorem

$\exists$ has a winning strategy in $\mathcal{E}(M, s, \varphi) \iff M, s \models \varphi$

To prove this theorem, it suffices to show:

1. If $M, s \models \varphi$, then $\exists$ has a *memoryless* winning strategy in the game $\mathcal{E}(M, s, \varphi)$.
2. If $M, s \not\models \varphi$, then $\forall$ has a *memoryless* winning strategy in the game $\mathcal{E}(M, s, \varphi)$.

A **memoryless strategy** is one in which the choice of next move depends only on the current position, not on the history of play.

In fact, $\exists$'s strategy always chooses positions (t, ψ) s.t. $M, t \models_{V'} \psi$, and similarly for $\forall$, where V' is a certain extension of V, which is obtained while fixpoint operators are unfolded during the game.

Constructing temporary valuations

We construct a temporary valuation V'.
List up the subformulas of φ with fixpoint operators in prefix and order them by size as:

$$\eta_1 X_1.\psi_1,\ \eta_2 X_2.\psi_2,\ \ldots,\ \eta_n X_n.\psi_n.$$

Here, if $\eta_i X_i.\psi_i$ is a subformula of $\eta_j X_j.\psi_j$, then $j \leq i$. In such a case, X_i may appear free in $\eta_j X_j.\psi_j$, but not vice versa.
Then we set

$$V_0 := V, \quad V_{i+1} := V_i \cup \left\{ X_{i+1} \mapsto ||\eta_{i+1} X_{i+1}.\psi_{i+1}||^M_{V_i} \right\}.$$

Then the truth values of the subformulas (with free variables) of φ are all determined by $V' = V_n$, and so a strategy chooses a position where $M, t \models_{V'} \psi$ holds.

But this doesn't guarantee a winning strategy, especially when the game proceeds infinitely. For such cases, we need to ensure that the largest subformula appearing infinitely often is headed by a ν-operator.

Transfinite inductive definitions

To resolve the above, we must ensure that repeated unfolding of μ-operators terminates in finite steps. Before explaining our trick, we first recall the transfinite inductive definitions.

Let $\Psi(S)(= ||\theta(X)||^{M}_{V(X):=S})$ be a monotone operator. Then, define an increasing sequence:

$$\Psi^0 := \varnothing, \ \Psi^{n+1} := \Psi(\Psi^n), \ \Psi^\omega := \bigcup_{n<\omega} \Psi^n.$$

However, Ψ^ω may not be a fixpoint. Further, we may have $\Psi^{\omega+1} := \Psi(\Psi^\omega) \supsetneq \Psi^\omega$. In general, you can define $\Psi^{\alpha+1} := \Psi(\Psi^\alpha)$ for any ordinal α, and for a limit ordinal λ,

$$\Psi^\lambda := \bigcup_{\alpha<\lambda} \Psi^\alpha.$$

Thus, there exists an ordinal $\bar{\alpha}$ such that $\Psi^{\bar{\alpha}}$ is a fixpoint, that is,

$$\Psi^{\bar{\alpha}} = \bigcap\{S \subseteq W : \Psi(S) \subseteq S\}.$$

Ordinal-valued strategy

From the list of fixpoint subformulas of φ, we extract the subformulas headed by μ:

$$\mu_1 X_1.\theta_1,\ \mu_2 X_2.\theta_2,\ \ldots,\ \mu_m X_m.\theta_m$$

For an ordinal sequence $r = (\alpha_1, \ldots, \alpha_m)$, we define V^r_n as follows:

$$V^r_0 := V, \quad V^r_{i+1} := V^r_i \cup \left\{ X_{i+1} \mapsto \begin{cases} ||\eta_{i+1} X_{i+1}.\psi_{i+1}||^M_{V^r_i} & \text{if } \eta_{i+1} = \nu \\ ||\mu^{\alpha_j}_j X_j.\theta_j||^M_{V^r_i} & \text{if } \eta_{i+1} X_{i+1} = \mu X_j \end{cases} \right\}$$

where $||\mu^\alpha X.\theta(X)||^M_V := \Psi^\alpha$ for $\Psi(S) = ||\theta(X)||^M_{V(X):=S}$.

There exists a minimal ordinal sequence r (in lexicographic order) such that $M, t \models_{V^r_n} \psi$. Denote such an r by $r^\mu(t, \psi)$.

Key properties of r^μ

For a valid move in the game,

- If ψ is headed by ν, then r^μ remains unchanged.
- If ψ is constructed by Boolean or modal operators, r^μ stays the same or decreases.
- If ψ is headed by a μ-operator, then r^μ **strictly decreases**. This is possible because for $s \in ||\mu^\alpha X.\theta(x)||_V^M$, there exists some $\beta < \alpha$ such that $s \in ||\mu^\beta X.\theta(x)||_V^M$.

Winning strategy for $\exists$: Always choose (t, ψ) such that:

$$M, t \models_{V_n^{r^\mu(t,\psi)}} \psi.$$

Similarly, if $M, s \not\models_V \varphi$, define $r^\nu(t, \psi)$, and $\forall$'s strategy is to choose (t, ψ) such that:

$$M, t \not\models_{V_n^{r^\nu(t,\psi)}} \psi.$$

Thus, the adequacy theorem is proved.

Miscellaneous

Tips. ν for infinity (always) μ for finiteness (eventually)

- $\nu X.p \wedge \Box_a X$ $\cdots\cdots\cdots$ p always holds along every a-path.
- $\nu X.p \wedge \Box_a \Box_a X$ $\cdots\cdots\cdots$ p holds at every even position along every a-path.
- $\nu X.q \vee (p \wedge \Box_a X)$ $\cdots\cdots\cdots$ p holds until q holds along every a-path.
- $\mu X.p \vee \Diamond_a X$ $\cdots\cdots\cdots$ p eventually holds on some a-path.
- $\mu X.p \vee \Box_a X$ $\cdots\cdots\cdots$ p eventually holds on every a-path.
- $\mu X.q \vee (p \wedge \Box_a X)$ $\cdots$ along every a-path, p holds until q holds and q eventually holds.
- $\mu X.\nu Y.(p \wedge \Box_a X) \vee (\neg p \wedge \Box_a Y)$ $\cdots\cdots$ p holds only finitely often on every a-path.
- $\nu X.\mu Y.(p \wedge \Diamond_a X) \vee \Diamond_a Y)$ $\cdots\cdots$ p holds infinitely often on some a-path.
- $\nu X.\mu Y.\Diamond_a X \vee \Diamond_b Y$ $\cdots\cdots\cdots$ there exists a $\{a, b\}$-path with infinitely many a.

Problems

Prob 2. For two $\mathcal{L}_\mu$-formulas φ, ψ, we write $\varphi \equiv \psi$ if for any model M, $[\![\varphi]\!]^M = [\![\psi]\!]^M$. Show the following equivalences.
(1) $\mu X.\varphi(X) \equiv \varphi(\mu X.\varphi(X))$. (2) $\mu X.\varphi(X) \equiv \neg\nu X.\neg\varphi(\neg X)$.

Prob 3. Find an example of a model M and a formula $\theta(X)$ such that $[\![\mu^\omega X.\theta(X)]\!]$ is not equivalent to the fixpoint $[\![\mu X.\theta(x)]\!]$. For such an example, consider which ordinal α makes $[\![\mu^\alpha X.\theta(x)]\!]$ and $[\![\mu X.\theta(x)]\!]$ equivalent.

§5.4. Applications: CTL

As applications of μ-calculus, we introduce CTL (Computation Tree Logic), a representative temporal logic, and two related systems. These logics are widely used in computer science.

Although CTL may have different modal operators, we only focus on two binary operators:

$$\mathsf{A}(\varphi \mathsf{U} \psi), \quad \mathsf{E}(\varphi \mathsf{U} \psi)$$

These mean:

- $\mathsf{A}(\varphi \mathsf{U} \psi)$: On *all* infinite paths, ψ holds eventually and φ holds until then.
- $\mathsf{E}(\varphi \mathsf{U} \psi)$: On *some* path, ψ holds eventually and φ holds until then.

Here, a path is a sequence of states generated by computation. U stands for "Until". We also include the usual modal operator $\Box$.

Syntax of CTL

Definition (CTL Formula)

Let p be an atomic proposition. A CTL formula is inductively defined as:

$$\varphi ::= p \mid \neg\varphi \mid (\varphi \to \varphi) \mid \Box\varphi \mid \mathsf{A}(\varphi \mathsf{U} \varphi) \mid \mathsf{E}(\varphi \mathsf{U} \varphi)$$

Other logical connectives are defined as usual.
Unless stated otherwise, the frame (W, R) used in CTL is assumed to be **serial**: $\forall s\, \exists t\, (sRt)$.

Semantics of CTL

Definition (Satisfaction Relation)

Let $M = (W, R, v)$ be a relational model. The satisfaction relation $M, s \models \varphi$ is defined inductively as follows.
The usual clauses and $\Box\varphi$ apply.

- For p, $\neg\varphi$, $\varphi \to \psi$, the usual clauses apply.
- $M, s_0 \models \mathsf{A}(\varphi \mathsf{U} \psi)$: For every infinite path $s_0 R s_1 R s_2 R \ldots$, there exists i such that $M, s_i \models \psi$ and $M, s_j \models \varphi$ for all $j < i$.
- $M, s_0 \models \mathsf{E}(\varphi \mathsf{U} \psi)$: There exists a finite path $s_0 R s_1 R \ldots R s_i$ such that $M, s_i \models \psi$ and $M, s_j \models \varphi$ for all $j < i$.

Derived Operators and μ-Calculus Translation

Various modal operators can be defined from the above. For example,

$$\mathsf{AG}\varphi := \neg\mathsf{E}(\top\mathsf{U}\neg\varphi), \quad \text{where } \top := p \to p$$

Importantly, CTL formulas can be translated into modal μ-calculus formulas. Let $(\varphi)^\natural$ be the translation of a CTL formula φ into a modal μ-calculus formula:

$$(\mathsf{A}(\varphi\mathsf{U}\psi))^\natural := \mu X.\psi^\natural \vee (\varphi^\natural \wedge \Box X)$$

$$(\mathsf{E}(\varphi\mathsf{U}\psi))^\natural := \mu X.\psi^\natural \vee (\varphi^\natural \wedge \Diamond X)$$

All other connectives remain unchanged. Note: only one variable X is used in this translation.

Theorem: Translation to Modal μ-Calculus

Theorem 1 (Translation Theorem)

For any serial relational structure M and any CTL formula φ:

$$M, s \models \varphi \iff M, s \models (\varphi)^\natural$$

A CTL formula φ is said to be **valid** if $M, s \models \varphi$ for every serial model $M = (W, R, v)$ and every state s, and we write:

$$\models_{\mathrm{CTL}} \varphi$$

Validity in CTL can be reduced to the validity of the single-variable modal μ-calculus.

Thank you for your attention!