11 Model theory of the modal μ-calculus

In this Chapter we will see how to apply the automata-theoretic tools developed in the previous chapter to prove some model-theoretic results about the modal μ-calculus.

▶ `overview of chapter to be supplied`

11.1 Small model property

As our first result we will prove a small model property for the modal μ-calculus, by showing that if a modal automaton accepts some pointed Kripke model, it accepts one of which the size is bounded by the size of the automaton. Recall that, given a modal automaton $\mathbb{A}$ we refer to the class of pointed Kripke models that are accepted by $\mathbb{A}$ as the *query* of $\mathbb{A}$, notation: $\mathcal{Q}(\mathbb{A})$, and that classes of this form are called *recognizable*.

Theorem 11.1 *Let $\mathbb{A}$ be a modal automaton. Then $\mathcal{Q}(\mathbb{A}) \neq \varnothing$ iff $\mathbb{A}$ accepts a finite pointed model of size at most exponential in the state-size of $\mathbb{A}$.*

Because of the Simulation Theorem it suffices to prove Theorem 11.1 for *disjunctive* modal automata. Our proof will be based on an alternative perspective of these devices, revealing their close resemblance the Kripke models that they operate on.

Kripke automata

The key observation in our proof is that the semantics of the cover modality and the notion of a bisimulation are defined in a very similar fashion, both involving the coalgebraic presentation of Kripke models, and the notion of *relation lifting.*

Fix a set P of proposition letters. Recall from Remark 1.3 and Definition 1.4 that we can represent a Kripke model[9] (S, R, V) as a pair

$$\mathbb{S} = (S, \sigma : S \to \mathsf{K}S),$$

where K is the Kripke functor given by putting, for an arbitrary set S:

$$\mathsf{K}S := \wp(\mathsf{P}) \times \wp(S).$$

In Definition 1.29 we introduced two notions of *relation lifting.* Given a binary relation $Z \subseteq S \times S'$, we define the relation $\overline{\wp}Z \subseteq \wp S \times \wp S'$ as follows:

$$\begin{array}{rcll} \overline{\wp}Z & := & \{(X, X') \mid & \text{for all } x \in X \text{ there is an } x' \in X' \text{ with } (x, x') \in Z \\ & & \& & \text{for all } x' \in X' \text{ there is an } x \in X \text{ with } (x, x') \in Z\}. \end{array}$$

Similarly, define, associated with the Kripke functor K, the relation $\overline{\mathsf{K}}Z \subseteq \mathsf{K}S \times \mathsf{K}S'$ as follows:

$$\overline{\mathsf{K}}Z := \{((\pi, X), (\pi', X')) \mid \pi = \pi' \text{ and } (X, X') \in \overline{\wp}Z\}.$$

[9] We restrict to the monomodal case in this section.

Position	Player	Admissible moves
$(a, s) \in A \times S$	-	$\{(\alpha(a), \sigma(s))\}$
$(\beta, \tau) \in \mathsf{K}A \times \mathsf{K}S$	$\exists$	$\{Z \in \wp(A \times S) \mid (\beta, \tau) \in \overline{\mathsf{K}}Z\}$
$Z \in \wp(A \times S)$	$\forall$	$Z = \{(b, t) \mid (b, t) \in Z\}$

Table 23: Bisimilarity game for Kripke models

To make our point we now introduce a new class of automata, consisting of so-called *Kripke automata*, and show that these are in fact equivalent to the disjunctive automata defined earlier on.

As our starting point we consider, for two Kripke models $\mathbb{A} = \langle A, \alpha \rangle$ and $\mathbb{S} = \langle S, \sigma \rangle$, the bisimilarity game $\mathcal{B}(\mathbb{A}, \mathbb{S})$ of Definition 1.26. Using the above notion of relation lifting, the rules of this game can be reformulated as in Table 23. Recall that the winning conditions of the bisimilarity game are such that all infinite games are won by $\exists$.

The main conceptual step is to think of $\mathbb{A}$ as a 'proto-automaton' that we use to *classify* $\mathbb{S}$ rather than as of a Kripke model that we are comparing with $\mathbb{S}$. In order to turn $\mathbb{A}$ into a proper Kripke automaton, four technical modifications have to be made:

(1) A small change is that we require $\mathbb{A}$ (i.e., its carrier set A) to be finite.

(2) Second, and equally undramatic, we add an initial state to the structure of $\mathbb{A}$.

(3) Third, whereas the winner of an infinite match of a bisimulation game is always $\exists$, the winner of an infinite acceptance match will be determined by an explicit acceptance condition on A^ω — a parity condition, in our case.

(4) The fourth and foremost modification is that we introduce *nondeterminism* to the transition structure of $\mathbb{A}$. That is, Kripke automata will harbour many 'realizations' of Kripke models — and in each round of the acceptance game, it is $\exists$'s task to pick an actual local realization of the current state of $\mathbb{A}$.

Definition 11.2 Given a set P of proposition letters, a *Kripke automaton* for P is a quadruple $\mathbb{A} = \langle A, \Delta, \Omega, a_I \rangle$ such that the transition function Δ is given as a map $\Delta : A \to \wp(\mathsf{K}A)$. The *acceptance game* $\mathcal{A}(\mathbb{A}, \mathbb{S})$ associated with a Kripke automaton $\mathbb{A} = \langle A, \Delta, \Omega, a_I \rangle$ and a Kripke structure $\mathbb{S}$ is given by Table 24. A pointed Kripke model $(\mathbb{S}, s)$ is *accepted* by $\mathbb{A}$ if the position

Position	Player	Admissible moves	Priority
$(a, s) \in A \times S$	$\exists$	$\{(\gamma, \sigma(s)) \in \mathsf{K}A \times \mathsf{K}S \mid \gamma \in \Delta(a)\}$	$\Omega(a)$
$(\gamma, \tau) \in \mathsf{K}A \times \mathsf{K}S$	$\exists$	$\{Z \subseteq A \times S \mid (\gamma, \tau) \in \overline{\mathsf{K}}Z\}$	0
$Z \in \wp(A \times S)$	$\forall$	Z	0

Table 24: Acceptance game for Kripke automata

(a_I, s) is a winning position for $\exists$ in the acceptance game. $\lhd$

For an informal description of the acceptance game $\mathcal{A}(\mathbb{A}, \mathbb{S})$, note that each round consists of exactly three moves, with interaction pattern $\exists\exists\forall$. At a basic position (a, s), the 'K-unfolding' $\sigma(s) \in \mathsf{K}S$ of s is fixed, but $\exists$ *chooses* the unfolding of a to be an arbitrary element

γ of $\Delta(a)$. After this move, the play arrives at a position of the form $(\gamma, s) \in \mathsf{K}A \times S$. The players now proceed as in the bisimilarity game for Kripke models. First $\exists$ chooses a 'local bisimulation' linking γ and $\sigma(s)$, that is, a relation $Z \subseteq A \times S$ such that $(\gamma, \sigma(s)) \in \overline{\mathsf{K}}Z$. Spelled out, this means that $\exists$ can only choose such a relation Z if γ is of the form $(c, B) \in \wp(\mathsf{P}) \times \wp(A)$ with $c = \sigma_V(s)$, and that Z has to satisfy the back and forth conditions, stating that for all $b \in B$ there is $t \in R[s]$ with bZt, and vice versa. The round ends with $\forall$ choosing an element (b, t) from Z, thus providing the next basic position of the match.

We will now show that Kripke automata are nothing but disjunctive automata in disguise, and vice versa.

Definition 11.3 First let $\mathbb{A} = \langle A, \Delta, \Omega, a_I \rangle$ be some Kripke automaton. We define its modal companion $\mathbb{A}^M$ as the disjunctive modal automaton $\mathbb{A}^M := \langle A, \Delta^M, \Omega, a_I \rangle$, where $\Delta^M : A \times \wp(\mathsf{P}) \to \mathtt{1DML}(A)$ is given by putting

$$\Delta^M(a, c) := \bigvee \{\nabla B \mid (c, B) \in \Delta(a)\}.$$

Conversely, let $\mathbb{D} = \langle D, \Theta, \Omega, d_I \rangle$ be a disjunctive modal automaton. Without loss of generality we may assume that the domain of Θ consists of formulas in the restricted format of Remark 10.14, that is, for every pair $(a, c) \in A \times \wp(\mathsf{P})$ there is a (possibly empty) index set $I_{a,c}$ such that

$$\Theta(a, c) = \bigvee \{\nabla B_i \mid i \in I_{a,c}\}.$$

We now define the transition map Δ_Θ by putting

$$\Delta_\Theta(a) := \{(c, B_i) \in \mathsf{K}A \mid c \in \wp(\mathsf{P}), i \in I_{a,c}\},$$

and define $\mathbb{D}^K := \langle D, \Delta_\Theta, \Omega, d_I \rangle$ and call this structure the *Kripke companion* of $\mathbb{D}$. $\lhd$

Remark 11.4 For a better understanding of the equivalence between disjunctive modal automata and Kripke models, it may be useful to take the following perspective. Given sets P (of proposition letters) and A of states, it is not hard to see that the collection of possible transition functions of disjunctive modal automata (in the restricted format of Remark 10.14 corresponds to the set

$$T_D := \big(A \times \wp(\mathsf{P})\big) \to \wp(\wp(A)),$$

while the set of possible transition maps of Kripke automata is given as the collection

$$T_K := A \to \wp\big(\wp(\mathsf{P}) \times \wp(A)\big).$$

Now recall that by 'currying' there is a bijective correspondence
(†) $(X \times Y) \to Z \;\cong\; X \to (Y \to Z)$
for any triple of sets X, Y and Z. Furthermore, for any set X there is a well-known bijective correspondence between the powerset $\wp(X)$ of X and the collection of functions from X to the two-element set $2 := \{0, 1\}$:
(‡) $\wp(X) \;\cong\; X \to 2$.

Using these observations it is straightforward to verify the following bijective correspondences between the sets T_D and T_K:

$$\begin{array}{ll} & (A \times \wp(\mathsf{P})) \to \wp\wp(A) \\ \cong (\ddagger) & (A \times \wp(\mathsf{P})) \to (\wp(A) \to 2)) \\ \cong (\dagger) & (A \times \wp(\mathsf{P}) \times \wp(A)) \to 2 \\ \cong (\dagger) & A \to \Big((\wp(\mathsf{P}) \times \wp(A)) \to 2\Big) \\ \cong (\ddagger) & A \to \wp(\wp(\mathsf{P}) \times \wp(A)) \end{array}$$

In fact, the translations given in Definition 11.3 can be obtained by computing the bijections between T_D and T_K, on the basis of those in (†) and (‡). ◁

Proposition 11.5 *(i) Let* $\mathbb{A} = \langle A, \Delta, \Omega, a_I \rangle$ *be a Kripke automaton. Then* $\mathbb{A} \equiv \mathbb{A}^M$.
(ii) Let $\mathbb{D} = \langle D, \Theta, \Omega, d_I \rangle$ *be a disjunctive modal automaton. Then* $\mathbb{D} \equiv \mathbb{D}^K$.

Proof. The proof of this proposition is straightforward. If we merge the two moves of $\exists$ in each round of the acceptance game for Kripke automata into one, we may in fact show that, for any Kripke model $\mathbb{S}$, the acceptance games $\mathcal{A}(\mathbb{A}^M, \mathbb{S})$ and $\mathcal{A}(\mathbb{A}, \mathbb{S})$ are *isomorphic*, and similarly for the acceptance games $\mathcal{A}(\mathbb{D}^K, \mathbb{S})$ and $\mathcal{A}(\mathbb{D}, \mathbb{S})$. QED

Small model property for Kripke automata

We will now prove the small model property for *Kripke automata*. This framework allows us to prove a result that is quite a bit stronger than just a small model theorem: we may show that, if $\mathbb{A}$ is a Kripke automaton recognizing a non-empty query, then $\mathcal{Q}\mathbb{A}$ contains a Kripke model that 'lives inside' or *inhabits* $\mathbb{A}$.

Definition 11.6 Let $\mathbb{A} = \langle A, \Theta, \Omega, a_I \rangle$ be a Kripke automaton. If S is a subset of A, and $\sigma : S \to \mathsf{K}S$ is such that $\sigma(s) \in \Delta(s)$ for all $s \in S$, then we say that the Kripke model $\mathbb{S} = \langle S, \sigma \rangle$ *inhabits* $\mathbb{A}$. When we use this terminology for a pointed Kripke model $(\mathbb{S}, s)$, we require in addition that $s = a_I$. ◁

The key tool in our proof of the small model property will be the following *satisfiability game* that we may associate with a Kripke automaton. Intuitively the reader may think of this game as the simultaneous projection on $\mathbb{A}$ of all acceptance games of $\mathbb{A}$, as should become clear from the proof of Theorem 11.8 below.

Definition 11.7 Let $\mathbb{A} = \langle A, \Delta, \Omega, a_I \rangle$ be a Kripke automaton. Then the *satisfiability game* $\mathcal{S}(\mathbb{A})$ is given by Table 25. The winning condition for infinite matches is defined using the priority map for game positions (see the table) as a parity condition. ◁

One last remark before we formulate and prove the main technical result of this section: the proof of this theorem involves a crucial application of the Positional Determinacy of parity games.

Position	Player	Admissible moves	Priority
$a \in A$	$\exists$	$\Delta(a)$	$\Omega(a)$
$(c, B) \in \mathsf{K}A$	$\forall$	B	0

Table 25: Satisfiability game for Kripke automata

Theorem 11.8 *The following are equivalent, for any Kripke automaton* $\mathbb{A} = \langle A, \Theta, \Omega, a_I \rangle$*:*
1) $\mathcal{Q}(\mathbb{A}) \neq \varnothing$*;*
2) $a_I \in \mathrm{Win}_\exists(\mathcal{S}(\mathbb{A}))$*;*
3) $\mathbb{A}$ *accepts a pointed model inhabiting* $\mathbb{A}$*.*

Proof. $\boxed{1 \Rightarrow 2}$ Suppose that $\mathbb{A}$ accepts some pointed model $(\mathbb{S}, s_0)$. Then by definition, $\exists$ has a winning strategy in the acceptance game $\mathcal{A}(\mathbb{A}, \mathbb{S})@(a_I, s_0)$. This strategy will be the basis of her winning strategy in the satisfiability game of $\mathbb{A}$.

Concretely, in $\mathcal{S}(\mathbb{A})@a_I$, $\exists$ will maintain the following condition. Put $a_0 = a_I$, and let

$$a_0(c_1, B_1)a_1(c_2, B_2) \dots a_k,$$

be an initial segment of an $\mathcal{S}(\mathbb{A})$-match (with $(c_{i+1}, B_{i+1}) \in \Theta(a_i)$ being the move of $\exists$ at position a_i, and $a_{i+1} \in B_{i+1}$ the next move of $\forall$). Then $\exists$ sees this match as the projection of a parallel match of $\mathcal{A}(\mathbb{A}, \mathbb{S})@(a_I, s_0)$ where she plays her winning strategy:

$$\begin{array}{ccccccccc}
(a_0, s_0) & ((c_1, B_1), s_0) & Z_1 & (a_1, s_1) & \dots & (a_k, s_k) & ((c_{k+1}, B_{k+1}), s_k) & Z_{k+1} & \dots \\
\Downarrow & \Downarrow & \Downarrow & \Downarrow & & \Downarrow & \Downarrow & \Downarrow & \\
a_0 & (c_1, B_1) & - & a_1 & \dots & a_k & (c_{k+1}, B_{k+1}) & - & \dots
\end{array}$$

The existence of such a parallel match is easily proved by an inductive argument, of which the base case is immediate by the shape (a_I versus (a_I, s_0)) of the initial game positions. Inductively assume that at stage k, the matches of $\mathcal{S}(\mathbb{A})$ and $\mathcal{A}(\mathbb{A}, \mathbb{S})$ have arrived at the positions a_k and (a_k, s_k) respectively. We will show that there is a way to continue both matches for one round in such a way that the next basic positions are of the form b and (b, t), respectively, for some $b \in A$ and $t \in S$, with the continuation in the acceptance game being guided by $\exists$'s winning strategy.

Suppose that $\exists$'s winning strategy in the acceptance game tells her to choose position $((c, B), \sigma(s_k))$, followed by the relation Z. Then at position a_k of $\mathcal{S}(\mathbb{A})$, we define her strategy to be such that she picks (c, B). Now suppose that in the match of $\mathcal{S}(\mathbb{A})$, $\forall$ chooses some element $b \in B$ as the next position. It follows by the assumption that $\exists$'s strategy is winning, that $(c, B) \in \Theta(a_k)$, $c = \sigma_V(s_k)$ and $(B, R[s_k]) \in \overline{\wp}(Z)$. Hence there must be an element $t \in R[s_k]$ such that $(b, t) \in Z$; in the acceptance game, she may look at a continuation of the match where $\forall$ picks the pair (b, t). In other words, we have proved that $\exists$ can maintain the parallel match for one more round.

Using this strategy in the satisfiability game will then guarantee her to win the match, since the associated sequence of $\mathbb{A}$-states is the same for both matches, and in the $\mathcal{A}(\mathbb{A}, \mathbb{S})$-match $\exists$ plays according to a strategy that was assumed to be winning.

$\boxed{2 \Rightarrow 3}$ Assume that $\exists$ has a winning strategy in the satisfiability game starting from the initial state a_I of $\mathbb{A}$. Let $S := \text{Win}_\exists(\mathcal{S}(\mathbb{A}))$ be the set of positions in A that are winning for $\exists$. The key point of the satisfiability game for Kripke automata is that $\mathcal{S}(\mathbb{A})$ is a parity game, and so we may without loss of generality assume that this strategy is *positional*, see Theorem **??**. In other words, we may represent it as a map $\sigma : S \to \mathsf{K}A$. We invite the reader to check that $\sigma(a) \in \mathsf{K}S$ for all $a \in S$. Now define $\mathbb{S}$ be the Kripke model $\langle S, \sigma \rangle$. The map $\sigma : S \to \mathsf{K}S$ then induces a binary relation $R \subseteq S \times S$ and a valuation $V : \mathsf{P} \to \wp(S)$, viz., the unique R and V such that $\sigma(s) = (R[s], \sigma_V(s))$. We claim that $\mathbb{A}$ accepts $(\mathbb{S}, a_I)$.

To see why this is the case, we will prove that (a_I, a_I) is a winning position in the acceptance game $\mathcal{A}(\mathbb{A}, \mathbb{S})$. The winning strategy that we may equip $\exists$ with in this game is in fact very simple:

- at position (a, s), pick $(\sigma(a), \sigma(s))$ as the next position if $a = s \in \text{Win}_\exists(\mathcal{S}(\mathbb{A}))$, and choose a random element otherwise;
- at position $((c, B), (c', B'))$, pick the relation $\{(b, b) \mid b \in B \cap B'\}$.

It can be proved that any match of the acceptance game in which $\exists$ uses this strategy, can be 'projected' onto a match of the satisfiability game in which she plays her winning strategy:

$$\begin{array}{cccccccc}
(a_I, a_I) & (\sigma(a_I), \sigma(a_I)) & \{(b, b) \mid b \in R[a_I]\} & (a_1, a_1) & (\sigma(a_1), \sigma(a_1)) & \ldots & (a_n, a_n) & \ldots \\
\Downarrow & \Downarrow & \Downarrow & \Downarrow & \Downarrow & & \Downarrow & \\
a_I & \sigma(a_I) & - & a_1 & \sigma(a_1) & \ldots & a_n & \ldots
\end{array}$$

Given the winning conditions of $\mathcal{A}(\mathbb{A}, \mathbb{S})$ and $\mathcal{S}(\mathbb{A})$ it is then immediate that the given strategy indeed guarantees that $\exists$ wins any match starting at position (a_I, a_I).

$\boxed{3 \Rightarrow 1}$ This implication is a direct consequence of the definitions. QED

11.2 Normal forms and decidability

In this section we will see two more corollaries of the results in the previous chapter.

Disjunctive normal form

As a first consequence, we now see that every formula of the modal μ-calculus can be brought into so-called *disjunctive normal form*. For the definition of the connectives used below we refer to Definition 1.37.

Definition 11.9 Given a sets P of proposition letters, the set of *disjunctive* modal μ-calculus formulas[10] over P is given by the following grammar:

$$\varphi ::= x \mid \bot \mid \top \mid \bigvee \Phi \mid \alpha \bullet \Phi \mid \mu x.\varphi \mid \nu x.\varphi$$

Here x is a variable *not* in P, α denotes a finite set of literals over P, and Φ is a finite collection of disjunctive formulas over P.

[10] In this section we work with a variation of the earlier definition: it will be convenient for us here to have $\top$ and $\bot$ as *primitive* formulas.

We let $\mu\mathtt{ML}_{\mathsf{D}}(\mathsf{P})$ denote the sentences of this language, that is, the disjunctive formulas φ such that $FV(\varphi) \subseteq \mathsf{P}$. $\lhd$

These formula are called disjunctive because the only admissible conjunctions are the special ones of the form $\alpha \bullet \Phi$, where α is a propositional formula (in fact, a conjunction of literals).

Theorem 11.10 *There is an effective algorithm that rewrites a modal fixpoint formula $\xi \in \mu\mathtt{ML}(\mathsf{P})$ into an equivalent disjunctive formula ξ^d of closure size at most exponential in $|\xi|$.*

- ▶ `proof (based on the results of the previous chapters) to be supplied.`
- ▶ `size issues to be addressed!`

Decidability

- ▶ `Intro`

Theorem 11.11 *There is an algorithm that decides in linear time (measured in dag-size) whether a given disjunctive formula ξ is satisfiable or not.*

Proof. It is easy to see that the proof of this proposition is a direct consequence of the following observations:

1. $\top$ is satisfiable;
2. $\bot$ is not satisfiable;
3. $\bigvee \Phi$ is satisfiable iff some $\varphi \in \Phi$ is satisfiable;
4. $\alpha \bullet \nabla\Phi$ is satisfiable iff both α and each $\varphi \in \Phi$ is satisfiable;
5. if $\mu x.\varphi$ is disjunctive, then it is satisfiable iff $\varphi[\bot/x]$ is satisfiable;
6. if $\nu x.\varphi$ is disjunctive, then it is satisfiable iff $\varphi[\top/x]$ is satisfiable.

The proof of these claims is left as an exercise for the reader. QED

Decidability of the satisfiability problem for modal fixpoint formulas is then an immediate consequence of the previous two results.

Corollary 11.12 *There is an algorithm that decides in elementary time whether a given modal fixpoint formula ξ is satisfiable or not.*

- ▶ `Corollary 11.12 does not provide the best complexity bound for the satisfiability problem for the` μ`-calculus, which can in fact be solved in (singly) exponential time.`

11.3 Uniform interpolation and bisimulation quantifiers

In this section we will prove that the modal μ-calculus enjoys the property of *uniform interpolation* by proving that we can express the so-called *bisimulation quantifiers* in the language.

Definition 11.13 Given two modal fixpoint formulas φ and ψ, we say that ψ is a *(local) consequence* of φ, notation: $\varphi \models \psi$, if $\mathbb{S}, s \Vdash \varphi$ implies $\mathbb{S}, s \Vdash \psi$, for every pointed Kripke model $(\mathbb{S}, s)$. ◁

A formalism has the *(Craig) interpolation property* if we can find an *interpolant* for every pair of formulas φ and ψ such that $\varphi \models \psi$. This interpolant is a formula θ such that $\varphi \models \theta$ and $\theta \models \psi$; but most importantly, the requirement on θ is that it may only use proposition letters that occur both in φ and ψ, or more precisely: $FV(\theta) \subseteq FV(\varphi) \cap FV(\psi)$.

▶ `why this is an important property`

Uniform interpolation is a very strong version of interpolation in which the interpolant θ does not depend on the particular shape of one of the formulas, but only on its *vocabulary* (set of free variables). More precisely, we define the following.

Definition 11.14 Let φ be a modal fixpoint formula, and $\mathsf{P} \subseteq FV(\varphi)$ be a set of variables. Then a *(right) uniform interpolant* of φ with respect to P is a formula θ with $FV(\theta) \subseteq \mathsf{P}$, such that

$$\varphi \models \psi \text{ iff } \theta \models \psi. \tag{121}$$

for all formulas ψ with $FV(\psi) \cap FV(\varphi) \subseteq \mathsf{P}$. ◁

In words, (121) states that θ has exactly the same consequences as φ, at least, if we restrict to formulas ψ such that all free variables shared by φ and ψ belong to P.

Remark 11.15 To justify the terminology 'uniform interpolant', take some formula ψ with $FV(\psi) \cap FV(\varphi) \subseteq \mathsf{P}$. We claim that

$$\varphi \models \psi \text{ implies } \varphi \models \theta \text{ and } \theta \models \psi \tag{122}$$

for any uniform interpolant θ of φ with respect to P.

To see this, suppose that $\varphi \models \psi$, and let θ be a uniform interpolant of φ with respect to P. Then we have $\theta \models \psi$ by (121), so it remains to show that $\varphi \models \theta$. But this follows immediately from the fact that by definition we have $FV(\theta) \cap FV(\varphi) \subseteq \mathsf{P}$, so that we may apply (121) to θ itself (and use that, obviously, $\theta \models \theta$). ◁

Remark 11.16 Dually, we could have introduced the notion of a *left* uniform interpolant for ψ, instead of a *right* interpolant for φ. A left interpolant for ψ, with respect to a set $\mathsf{P} \subseteq FV(\psi)$ of proposition letters, is a formula χ with $FV(\chi) \subseteq \mathsf{P}$, and such that $\varphi \models \psi$ iff $\varphi \models \chi$. But since negation is definable in the modal μ-calculus as an operation $\sim : \mu\mathtt{ML}(\mathsf{P}) \to \mu\mathtt{ML}(\mathsf{P})$ and so we have $\varphi \models \psi$ iff $\sim\psi \models \sim\varphi$, it is not hard to see that if θ is a (right) uniform interpolant for ψ, then its negation $\sim\theta$ is a left interpolant for ψ. In other words, since our language is closed under classical negation, requiring that every formula has a right uniform interpolant is equivalent to requiring that every formula has a left uniform interpolant. ◁

The following theorem states that uniform interpolants exist in the modal μ-calculus.

Theorem 11.17 (Uniform Interpolation) *Let φ be a modal fixpoint formula, and let P be a set of variables such that $\mathsf{P} \subseteq FV(\varphi)$. Then φ has a uniform interpolant with repect to P.*

The proof consists of showing that the modal μ-calculus can express the so-called *bisimulation quantifiers.*

Definition 11.18 Given a proposition letter q, the *bisimulation quantifier* $\widetilde{\exists} q$ is an operator with the following semantics:

$$\mathbb{S}, s \Vdash \widetilde{\exists} q.\varphi \text{ iff } \mathbb{S}', s' \Vdash \varphi, \text{ for some pointed model } \mathbb{S}', s' \leftrightarrows_{\mathsf{R}\setminus q} \mathbb{S}, s, \tag{123}$$

where $\mathbb{S}$ is some Kripke model over a set R of proposition letters, and $\leftrightarrows_{\mathsf{R}\setminus q}$ is the bisimilarity relation 'up to q', that is, we only require the condition (prop) of Definition 1.19 to hold for proposition letters $p \in \mathsf{R} \setminus q$. $\lhd$

The bisimulation quantifier $\widetilde{\exists} q$ is a second-order existential quantifier, but nonstandard in the sense that it does not quantify over subsets of the actual model $\mathbb{S}$, but rather over subsets of possibly distinct (but bisimilar-up-to-q) models. For instance, if s is a state in $\mathbb{S}$ with one single successor, then obviously the formula $\widetilde{\exists} q(\Diamond q \land \Diamond \overline{q})$ would be false if we had to interpret q as a subset of S. However, taking a bisimilar pointed model $(\mathbb{S}', s')$ such that s' has two successors, we can easily interpret q as a subset of S' such that the formula $\Diamond q \land \Diamond \overline{q}$ becomes true at s'. Similarly, the formula $\widetilde{\exists} q(q \land \Box \overline{q})$ holds at any point in any Kripke model.

The main result underlying the proof of Theorem 11.17 is that the bisimulation quantifiers are *definable* in the modal μ-calculus. The following notation will be convenient.

Convention 11.19 Where P is a set of proposition letters, and q is a proposition letter (which may or may not belong to P), we write $\mathsf{P} \setminus q$ rather than $\mathsf{P} \setminus \{q\}$.

Theorem 11.20 *For any set P of proposition letters, and any proposition letter q, there is a map*

$$\widetilde{\exists} q : \mu\mathtt{ML}_{\mathsf{D}}(\mathsf{P}) \to \mu\mathtt{ML}_{\mathsf{D}}(\mathsf{P} \setminus q)$$

such that for any formula $\varphi \in \mu\mathtt{ML}_{\mathsf{D}}(\mathsf{P})$, we have $FV(\widetilde{\exists} q.\varphi) = FV(\varphi) \setminus q$, and the semantics of $\widetilde{\exists} q.\varphi$ satisfies (123)*, for any Kripke model over a set of proposition letters $\mathsf{R} \supseteq \mathsf{P}$.*

The proof of Theorem 11.20 crucially involves *disjunctive* modal automata. Before going into the details, there is a technicality that we need to get out of the way.

Remark 11.21 Let $\mathbb{A} = \langle A, \Theta, \Omega, a_I \rangle$ be a modal automaton over some set P of proposition letters, and let $\mathbb{S} = (S, R, V)$ be a Kripke model over some, possibly larger, set R. Then strictly speaking the acceptance game $\mathcal{A}(\mathbb{A}, \mathbb{S})$ is not well-defined since the domain of the transition map Θ is of the form $\mathsf{Dom}(\Theta) = A \times \wp(\mathsf{P})$, while the range of the colouring map

σ_V of $\mathbb{S}$ is the set $\mathsf{Ran}(\sigma_V) = \wp(\mathsf{R})$. But clearly we can take care of this mismatch by working with the map $\Theta_{\mathsf{R}} : A \times \wp(\mathsf{R}) \to \mathtt{1ML}(A)$ given by

$$\Theta_{\mathsf{R}}(a, c) := \Theta(a, c \cap \mathsf{P}).$$

In the sequel we will largely ignore this issue. $\lhd$

We now turn to the details of the proof of Theorem 11.20. Because of the existence of truth-preserving translations between formulas and automata, it suffices to provide a construction on modal automata that instantiates the bisimulation quantifier, and because of the Simulation Theorem it suffices to define this construction for disjunctive modal automata.

Definition 11.22 Let P be a set of proposition letters and let q be a proposition letter (possibly but not necessarily in P). Let $\mathbb{A} = \langle A, \Theta, \Omega, a_I \rangle$ be a disjunctive modal automaton over the set P. We abbreviate $C := \wp(\mathsf{P})$ and $C^- := \wp(\mathsf{P} \setminus \{q\})$.

Now we define the modal automaton $\widetilde{\exists} q.\mathbb{A}$ as the structure $\widetilde{\exists} q.\mathbb{A} := \langle A, \Theta^{\pm q}, \Omega, a_I \rangle$, where

$$\Theta^{\pm q}(a, c) := \Theta(a, c \setminus \{q\}) \vee \Theta(a, c \cup \{q\})$$

defines the transition map $\Theta^{\pm q} : A \times C^- \to \mathtt{1DML}(A)$. $\lhd$

The main technical result that we will prove is the following. Recall from Definition 10.16 that we write $\mathbb{S}, s_I \Vdash_s \mathbb{A}$ in case $\exists$ has a functional strategy in the game $\mathcal{A}(\mathbb{A}, \mathbb{S})@(a_I, s_I)$.

Proposition 11.23 *Let $\mathbb{A}$ be a disjunctive modal P-automaton, and let $\mathbb{S}$ be a Kripke model over some set $\mathsf{R} \supseteq \mathsf{P}$. Then the following are equivalent, for any state $s_I \in \mathbb{S}$:*
1) $\mathbb{S}, s_I \Vdash_s \widetilde{\exists} q.\mathbb{A}$;
2) $\mathbb{S}[q \mapsto Q], s_I \Vdash_s \mathbb{A}$, for some subset $Q \subseteq S$.

Proof. We only consider the case where $\mathsf{R} = \mathsf{P}$, leaving it for the reader to extend the result to the more general case (cf. Remark 11.21). Fix a disjunctive P-automaton $\mathbb{A} = \langle A, \Theta, \Omega, a_I \rangle$ and an R-model $\mathbb{S} = (S, R, V)$; to simplify notation we will write $c_t := \sigma_V(t)$, for an arbitrary point $t \in S$. Similarly, we will write $c - q := c \setminus \{q\}$ and $c + q := c \cup \{q\}$ for an arbitrary colour $c \in \wp(\mathsf{P})$. Furthermore, we will use the one-step presentation of the acceptance game, as in Table 21.

For the direction 1) $\Rightarrow$ 2) of the Proposition, assume that $\mathbb{S}, s_I \Vdash_s \widetilde{\exists} q.\mathbb{A}$. In other words, $\exists$ has a *functional* positional strategy f which is winning in the game $\mathcal{A}(\widetilde{\exists} q.\mathbb{A}, \mathbb{S})@(a_I, s_I)$. Abbreviate $\mathcal{A} := \mathcal{A}(\widetilde{\exists} q.\mathbb{A}, \mathbb{S})$.

Let $U \subseteq S$ be the set of points t in S such that, for some state $a \in A$, the position (a, t) is f-reachable in $\mathcal{A}@(a_I, s_I)$. It follows from functionality of f that for every $t \in U$ there is a *unique* such state in $\mathbb{A}$; we will denote this state as a_t. Furthermore, since f is a *winning strategy* in $\mathcal{A}@(a_I, s_I)$, every position of the form (a_t, t) is winning for $\exists$, and so by legitimacy of f, the marking $m_t : R[t] \to \wp(A)$ picked by f at this position is such that

$$(R[t], m_t) \Vdash^1 \Theta^{\pm q}(a_t, c_t). \tag{124}$$

Given that $\Theta^{\pm q}(a_t, c_t) = \Theta(a_t, c_t - q) \vee \Theta(a_t, c_t + q)$, this observation provides the set $Q \subseteq S$ that we are looking for:

$$Q := \{t \in U \mid (R[t], m_t) \Vdash^1 \Theta(a_t, c_t + q)\}.$$

We claim that $\mathbb{S}[q \mapsto Q], s_I \Vdash_s \mathbb{A}$, and to show this, we define the following positional strategy f_Q for $\exists$ in $\mathcal{A}_Q := \mathcal{A}(\mathbb{A}, \mathbb{S}[q \mapsto Q])$. At a position $(a, t) \in A \times S$, $\exists$ will play as follows:

- in case $t \in U$ and $a = a_t$, she picks the marking m_t;
- in all other cases she picks a random marking.

We first show that for each $t \in U$ and $a = a_t$ this strategy provides a legitimate move in $\mathcal{A}_Q$, that is,

$$(R[t], m_t) \Vdash^1 \Theta(a_t, \sigma_{V[q \mapsto Q]}(t)). \tag{125}$$

To see this, make the following case distinction:

- If $(R[t], m_t) \Vdash^1 \Theta(a_t, c_t + q)$ then by definition of Q we find $t \in Q$. This means that $\sigma_{V[q \mapsto Q]}(t) = \sigma_V(t) \cup \{q\} = c_t + q$. In other words, (125) holds indeed.
- If, on the other hand, $(R[t], m_t) \not\Vdash^1 \Theta(a_t, c_t + q)$ then by definition of Q we find $t \notin Q$. Furthermore, by (124) and the definition of $\Theta^{\pm q}$ it must be the case that $(R[t], m_t) \Vdash^1 \Theta(a_t, c_t - q)$. But since $t \notin Q$ we have $\sigma_{V[q \mapsto Q]}(t) = \sigma_V(t) \setminus \{q\} = c_t - q$, so that again we obtain (125).

It remains to show that f_Q is functional, and *winning* for $\exists$ in $\mathcal{A}_Q@(a_I, s_I)$, but this is in fact easy. The point is that at any position of the form (a_t, t) the strategies f and f_Q prescribe the *same* move, viz., m_t, and that at the position m_t the moves of $\forall$ in $\mathcal{A}$ and $\mathcal{A}_Q$ are the same. From this it follows that every position for $\exists$ that is reachable in an f_Q-guided match of $\mathcal{A}_Q@(a_I, s_I)$ is of the form (a_t, t) (with $t \in U$), and so by our previous claim about the legitimacy of f_Q at such positions, f_Q is a surviving strategy. Now consider an f_Q-guided full match of $\mathcal{A}_Q@(a_I, s_I)$; this very same match is also an f-guided match of $\mathcal{A}$, and hence won by $\exists$ — after all we assumed that f is a winning strategy for $\exists$ in $\mathcal{A}(a_I, s_I)@(a_I, s_I)$, and the winning conditions in $\mathcal{A}_Q$ and $\mathcal{A}$ are the same. In other words, every f_Q-guided full match of $\mathcal{A}_Q@(a_I, s_I)$ is won by $\exists$. Finally, since f is a functional strategy, so is f_Q. This finishes the proof that 1) $\Rightarrow$ 2).

The proof of the opposite implication, 2) $\Rightarrow$ 1), is similar; we omit the details. QED

From this, Theorem 11.20 is almost immediate.

Proof of Theorem 11.20. Let P and q be a set of proposition letters and a proposition letter, respectively, let $\mathbb{A}$ be a disjunctive modal automaton over P, and let $(\mathbb{S}, r)$ be a pointed model over a set R of proposition letters such that $\mathsf{P} \subseteq \mathsf{R}$. It suffices to show that

$$\mathbb{S}, r \Vdash \widetilde{\exists} q.\mathbb{A} \text{ iff } \mathbb{S}', r' \Vdash \mathbb{A}, \text{ for some } (\mathbb{S}', r') \text{ with } \mathbb{S}, r \leftrightarroweq_{\mathsf{R} \setminus q} \mathbb{S}', r'. \tag{126}$$

But since $\mathbb{A}$ is disjunctive, it is easy to see that $\widetilde{\exists} q.\mathbb{A}$ is disjunctive as well, and so it follows from Theorem 10.18 that

$$\mathbb{S}, r \Vdash \widetilde{\exists} q.\mathbb{A} \text{ iff } \mathbb{S}', r' \Vdash_s \widetilde{\exists} q.\mathbb{A}, \text{ for some } (\mathbb{S}', r') \text{ with } \mathbb{S}, r \leftrightarroweq_{\mathsf{R}\setminus q} \mathbb{S}', r'. \tag{127}$$

Combining this with Proposition 11.23 we find

$$\mathbb{S}, r \Vdash \widetilde{\exists} q.\mathbb{A} \text{ iff } \mathbb{S}'[q \mapsto Q], r' \Vdash_s \mathbb{A}, \text{ for some } (\mathbb{S}', r') \text{ with } \mathbb{S}, r \leftrightarroweq_{\mathsf{R}\setminus q} \mathbb{S}', r' \text{ and some } Q \subseteq S'. \tag{128}$$

Now it is obvious that $\mathbb{S}'[q \mapsto Q], r' \leftrightarroweq_{\mathsf{R}\setminus q} \mathbb{S}', r'$. But then (126) is immediate. QED

Finishing this section, we show how to derive the uniform interpolation property from the definability of the bisimulation quantifiers.

Proof of Theorem 11.17. Fix the formula φ and the set P, and let $q_1, \ldots, q_n$ enumerate the free variables of φ that are *not* in P, that is, $\{q_1, \ldots, q_n\} = FV(\varphi) \setminus \mathsf{P}$. We claim that the formula $\widetilde{\exists} q_1 \cdots \widetilde{\exists} q_n.\varphi$ is the required (right) uniform interpolant of φ with respect to P.

To prove this, take an arbitrary formula ψ such that $FV(\psi) \cap FV(\varphi) \subseteq \mathsf{P}$. Clearly this implies that no q_i is a free variable of ψ. We first show that

$$\varphi \models \widetilde{\exists} q_1 \cdots \widetilde{\exists} q_n.\varphi.$$

To see this, let $(\mathbb{S}, s)$ be some pointed Kripke model (over some set $\mathsf{R} \supseteq FV(\varphi)$) such that $\mathbb{S}, s \Vdash \varphi$. Since we obviously have that $\mathbb{S}, s \leftrightarroweq_{\mathsf{R}\setminus q} \mathbb{S}, s$ for *any* proposition letter q, it easily follows that $\varphi \models \widetilde{\exists} q_1 \cdots \widetilde{\exists} q_n.\varphi$. This takes care of the right-to-left direction from (121).

For the opposite direction of (121), assume that $\varphi \models \psi$, and let $(\mathbb{S}, s)$ be a pointed Kripke model such that $\mathbb{S}, s \Vdash \widetilde{\exists} q_1 \cdots \widetilde{\exists} q_n.\varphi$. It follows that there is a sequence $(\mathbb{S}_i, s_i)_{0 \leq i \leq n}$ of pointed models such that $(\mathbb{S}, s) = (\mathbb{S}_0, s_0)$, $\mathbb{S}_n, s_n \Vdash \varphi$, and $\mathbb{S}_i, s_i \leftrightarroweq_{\mathsf{R}\setminus q_{i+1}} \mathbb{S}_{i+1}, s_{i+1}$ for all i with $0 \leq i < n$. Then by assumption it follows from $\mathbb{S}_n, s_n \Vdash \varphi$ that $\mathbb{S}_n, s_n \Vdash \psi$. But since none of the proposition letters q_i is free in ψ, step by step applying the bisimulation invariance of the modal μ-calculus we may show that each pointed model $\mathbb{S}_i, s_i$ satisfies ψ. In particular, we find that $\mathbb{S}, s \Vdash \psi$, as required. QED

Notes

The decidability of the satisfiability problem of the modal μ-calculus was first proved by Kozen and Parikh [11] via a reduction to SnS. Emerson & Jutla [?] established the EXPTIME-completeness of this problem. The finite model property was proved by Kozen [?].

Uniform interpolation of the modal μ-calculus was proved by D'Agostino & Hollenberg [?], who established some other model-theoretic results as well.

Exercises

Exercise 11.1 Let γ be some disjunctive fixed point formula.

(a) Show that $\mu x.\gamma$ is satisfiable iff $\gamma[\bot/x]$ is satisfiable.

(b) Show that $\nu x.\gamma$ is satisfiable iff $\gamma[\top/x]$ is satisfiable.

(c) Do the above statements hold for arbitrary fixed point formulas as well?

Exercise 11.2 Prove the left-to-right direction of (126) in Proposition 12.28.

Exercise 11.3 Is disjunctivity of the automaton $\mathbb{A}$ needed in the proof of Proposition 11.23?

Exercise 11.4 (PDL + bisimulation quantifier) Consider a setting with finitely many atomic actions. Let PDL+$\widetilde{\exists}$ be the extension of propositional dynamic logic with (explicit) bisimulation quantifiers. Show that there is a (truth-preserving) translation from the modal μ-calculus to PDL+$\widetilde{\exists}$.