14 Expressive Power of Monadic Second-Order Logic and Modal μ-Calculus

Philipp Rohde

Lehrstuhl für Informatik VII
RWTH Aachen

14.1 Introduction

We consider monadic second order logic (MSO) and the modal μ-calculus (L_μ) over transition systems (Kripke structures). It is well known that every class of transition systems which is definable by a sentence of L_μ is definable by a sentence of MSO as well. It will be shown that the converse is also true for an important fragment of MSO: every class of transition systems which is MSO-definable and which is closed under bisimulation – i.e., the sentence does not distinguish between bisimilar models – is also L_μ-definable. Hence we obtain the following expressive completeness result: the bisimulation invariant fragment of MSO and L_μ are equivalent. The result was proved by David Janin and Igor Walukiewicz. Our presentation is based on their article [91]. The main step is the development of automata-based characterizations of L_μ over arbitrary transition systems and of MSO over transition trees (see also Chapter 16). It turns out that there is a general notion of automaton subsuming both characterizations, so we obtain a common ground to compare these two logics. Moreover we need the notion of the ω-unravelling for a transition system, on the one hand to obtain a bisimilar transition tree and on the other hand to increase the possibilities of choosing successors.

We start with a section introducing the notions of transition systems and transition trees, bisimulations and the ω-unravelling. In Section 14.3 we repeat the definitions of MSO and L_μ. In Section 14.4 we develop a general notion of automaton and acceptance conditions in terms of games to obtain the characterizations of the two logics. In the last section we will prove the main result mentioned above.

14.2 Preliminary Definitions

Let $\text{Prop} = \{p, p', \dots\} \cup \{\bot, \top\}$ be a set of unary predicate symbols (propositional letters) and $\text{Rel} = \{r, r', \dots\}$ a set of binary predicate symbols (letters for relations). We consider a signature containing only symbols from Prop and Rel. Let $\text{Var} = \{X, Y, \dots\}$ be a countable set of variables.

Definition 14.1. Let $S^{\mathcal{M}}$ be a non-empty set of states and $\text{sr}^{\mathcal{M}}$ an element of $S^{\mathcal{M}}$. For each $r \in \text{Rel}$ let $r^{\mathcal{M}}$ a binary relation on $S^{\mathcal{M}}$ and for each $p \in \text{Prop}$

E. Grädel et al. (Eds.): Automata, Logics, and Infinite Games, LNCS 2500, pp. 239-257, 2002.

let $p^{\mathcal{M}} \subseteq S^{\mathcal{M}}$ a subset of $S^{\mathcal{M}}$. A **transition system $\mathcal{M}$ with source sr$^{\mathcal{M}}$** – transition system for short – is the tuple

$$\big(S^{\mathcal{M}}, \mathrm{sr}^{\mathcal{M}}, \{\, r^{\mathcal{M}} \mid r \in \mathrm{Rel}\,\}, \{\, p^{\mathcal{M}} \mid p \in \mathrm{Prop}\,\}\big)\,.$$

For every $r \in \mathrm{Rel}$ and state $s \in S^{\mathcal{M}}$ let

$$\mathrm{scc}_r^{\mathcal{M}}(s) := \{\, s' \in S^{\mathcal{M}} \mid (s, s') \in r^{\mathcal{M}}\,\}$$

be the set of r-successors of s.

A transition system $\mathcal{M}$ is called a **transition tree** if for every state $s \in S^{\mathcal{M}}$ there is a unique path to the root of the tree (alias the source of the system), i.e., a unique finite sequence $s_0, \ldots, s_n$ in $S^{\mathcal{M}}$ such that $s_0 = \mathrm{sr}^{\mathcal{M}}$, $s_n = s$ and for every $i \in \{0, \ldots, n-1\}$ we have $s_{i+1} \in \mathrm{scc}_{r_i}^{\mathcal{M}}(s_i)$ for exactly one $r_i \in \mathrm{Rel}$.

Definition 14.2. Two transition systems $\mathcal{M}$ and $\mathcal{N}$ are **bisimilar** – denoted by $\mathcal{M} \sim \mathcal{N}$ – if there is a **bisimulation relation** $R \subseteq S^{\mathcal{M}} \times S^{\mathcal{N}}$ such that for every $(s, t) \in R$, $p \in \mathrm{Prop}$ and $r \in \mathrm{Rel}$:

- $\big(\mathrm{sr}^{\mathcal{M}}, \mathrm{sr}^{\mathcal{N}}\big) \in R$,
- s satisfies p in $\mathcal{M}$ iff t satisfies p in $\mathcal{N}$, i.e., $s \in p^{\mathcal{M}} \iff t \in p^{\mathcal{N}}$ holds,
- "Zig": for every r-successor s' of s in $\mathcal{M}$ there exists a r-successor t' of t in $\mathcal{N}$ such that $(s', t') \in R$,
- "Zag": for every r-successor t' of t in $\mathcal{N}$ there exists a r-successor s' of s in $\mathcal{M}$ such that $(s', t') \in R$.

Bisimulations are also known as *zigzagrelations*. For an example of a bisimulation see Fig. 14.1.

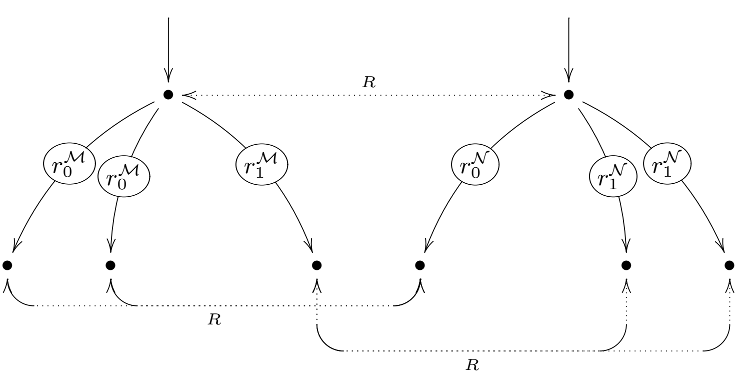

Fig. 14.1. Two bisimilar transition systems $\mathcal{M}$ and $\mathcal{N}$.

Let $\mathcal{C}$ be a class of transition systems. We say $\mathcal{C}$ is **bisimulation closed** if for all transition systems $\mathcal{M}$ and $\mathcal{N}$ the following holds:

$$\mathcal{M} \in \mathcal{C} \wedge \mathcal{N} \sim \mathcal{M} \Longrightarrow \mathcal{N} \in \mathcal{C}.$$

Exercise 14.1. Show that $\sim$ is an equivalence relation on the class of all transition systems.

Definition 14.3. Let $\mathcal{M}$ be a transition system and $s \in S^{\mathcal{M}}$. An ***ω*-path to *s*** is a finite sequence

$$s_0(a_1, r_1, s_1)(a_2, r_2, s_2) \ldots (a_n, r_n, s_n),$$

where $s_0 = \mathrm{sr}^{\mathcal{M}}$, $s_n = s$, $a_i \in \omega$ and each s_{i+1} is a r_i-successor of s_i, i.e., $s_{i+1} \in \mathrm{scc}_{r_i}^{\mathcal{M}}(s_i)$ holds for every $i \in \{0, \ldots, n-1\}$. The ***ω*-unravelling $\widehat{\mathcal{M}}$** of $\mathcal{M}$ is the transition system defined as follows:

- $S^{\widehat{\mathcal{M}}}$ is the set of the ω-paths to the elements of $S^{\mathcal{M}}$,
- the sources are identical: $\mathrm{sr}^{\widehat{\mathcal{M}}} = \mathrm{sr}^{\mathcal{M}}$,
- for $u, v \in S^{\widehat{\mathcal{M}}}$ and $r \in \mathrm{Rel}$ we set $(u, v) \in r^{\widehat{\mathcal{M}}}$ iff v is a one-term extension of u, i.e., there are $a \in \omega$ and $s \in S^{\mathcal{M}}$ such that $v = u(a, r, s)$,
- for $v \in S^{\widehat{\mathcal{M}}}$ and $p \in \mathrm{Prop}$ we set $v \in p^{\widehat{\mathcal{M}}}$ iff either $v = \mathrm{sr}^{\widehat{\mathcal{M}}}$ and $\mathrm{sr}^{\mathcal{M}} \in p^{\mathcal{M}}$ or $v = u(a, r, s)$ for some $u \in S^{\widehat{\mathcal{M}}}$, $a \in \omega$, $r \in \mathrm{Rel}$ such that $s \in p^{\mathcal{M}}$.

For an example of an ω-unravelling see Fig. 14.2.

Exercise 14.2. Let $\mathcal{M}$ be a transition system. Show that:

(1) The ω-unravelling $\widehat{\mathcal{M}}$ is always unique and a transition tree,
(2) $\mathcal{M}$ and $\widehat{\mathcal{M}}$ are bisimilar.
Hint: Consider the following relation $R \subseteq S^{\mathcal{M}} \times S^{\widehat{\mathcal{M}}}$:

$$(s, t) \in R :\Longleftrightarrow t \text{ is an } \omega\text{-path to } s.$$

The main property of the ω-unravelling for our purpose is, that we always have enough possibilities to choose a different r-successor for finitely many r-successors of an element in $\widehat{\mathcal{M}}$. In other words: Let t be an r-successor of s in $\mathcal{M}$ and let u be an ω-path to s. Then there are infinitely many r-successors of u in $\widehat{\mathcal{M}}$ which are bisimilar to t.

Definition 14.4. Let $\mathcal{M}$ and $\mathcal{N}$ be two transition systems. $\mathcal{M}$ is an **extension** of $\mathcal{N}$ – denoted by $\mathcal{M} \succeq \mathcal{N}$ – if there is a partial function $h : S^{\mathcal{M}} \to S^{\mathcal{N}}$ such that for every $s \in S^{\mathcal{M}}$, $r \in \mathrm{Rel}$ and $p \in \mathrm{Prop}$:

- the source in $\mathcal{M}$ is mapped to the source in $\mathcal{N}$: $h(\mathrm{sr}^{\mathcal{M}}) = \mathrm{sr}^{\mathcal{N}}$,
- s satisfies p in $\mathcal{M}$ iff $h(s)$ satisfies p in $\mathcal{N}$: $s \in p^{\mathcal{M}} \Longleftrightarrow h(s) \in p^{\mathcal{N}}$,
- s' is a r-successor of s in $\mathcal{M}$ if and only if $h(s')$ is a r-successor of $h(s)$ in $\mathcal{N}$: $h\big[\mathrm{scc}_r^{\mathcal{M}}(s)\big] = \mathrm{scc}_r^{\mathcal{N}}(h(s))$.

Exercise 14.3. Show that if $\mathcal{M}$ is an extension of $\mathcal{N}$ then $\mathcal{M}$ is bisimilar to $\mathcal{N}$, so the notion of bisimulation is more general than the notion of extension.

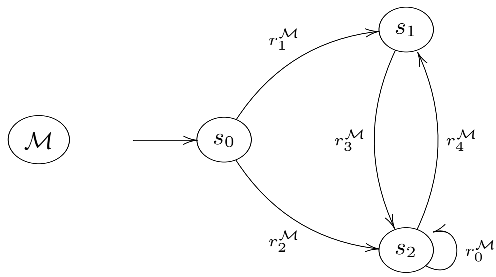

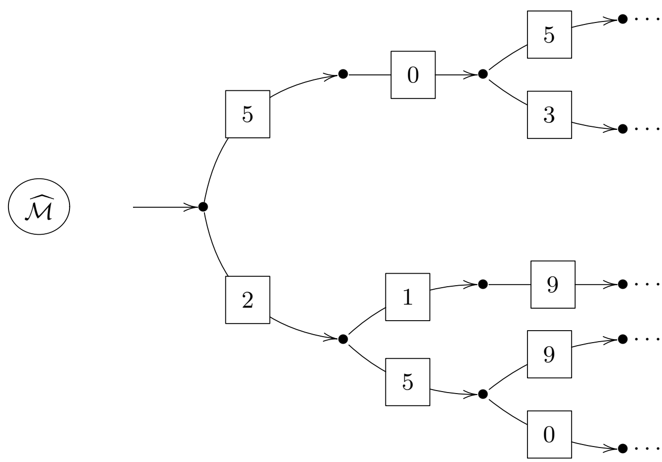

Fig. 14.2. A transition system $\mathcal{M}$ with source $s_0 = \mathrm{sr}^{\mathcal{M}}$ and a part of its ω-unravelling $\widehat{\mathcal{M}}$ (we suppressed the labelling of the nodes). Notice that in fact every node has infinitely many sons.

In fact we have:

Theorem 14.5 (Castellani 1987)**.** *Two transition systems $\mathcal{M}_1$ and $\mathcal{M}_2$ are bisimilar iff there is a transition system $\mathcal{N}$ such that $\mathcal{M}_1 \succeq \mathcal{N}$ and $\mathcal{M}_2 \succeq \mathcal{N}$.*

The proof can be found in [27]. Notice that one direction is the statement of the last exercise. $\mathcal{N}$ can be seen as quotient of $\mathcal{M}_1$ and $\mathcal{M}_2$ under bisimulation relation, i.e., the minimal representative of the equivalence class $[\mathcal{M}_1]_\sim$.

In the countable case we obtain:

Exercise 14.4. Let $\mathcal{M}$ and $\mathcal{N}$ be transition systems such that $S^{\mathcal{M}}$ and $S^{\mathcal{N}}$ are countable. Show that:

- $\widehat{\mathcal{M}}$ is an extension of $\mathcal{M}$,
- If $\mathcal{M}$ and $\mathcal{N}$ are bisimilar then $\widehat{\mathcal{M}}$ and $\widehat{\mathcal{N}}$ are isomorphic.

14.3 Monadic Second Order Logic and the Modal μ-Calculus

These two logics will be interpreted over transition systems. There are several ways to define MSO over transition systems, for example by using two types of variables (first-order and second-order variables) or by introducing a new predicate $\mathrm{sing}(X)$ for singleton sets. We use the following definition:

Definition 14.6. The signature of **monadic second order logic (MSO)** over transition systems contains unary predicate symbols from Prop, binary predicate symbols from Rel, the constant symbol sr and variables from Var. Formulae of MSO are defined inductively by the following grammar. Let $p \in \mathrm{Prop}$, $r \in \mathrm{Rel}$ and $X, Y \in \mathrm{Var}$:

- $\mathrm{sr}(X)$,
- $p(X)$,
- $r(X, Y)$,
- $X \subseteq Y$,
- $\neg\varphi$ for any formula φ,
- $\varphi \vee \psi$ for any formulae φ and ψ,
- $\exists X.\varphi(X)$ for any formula φ.

Other connectives, such as the conjunction $\wedge$, the implication $\Longrightarrow$ and the universal quantification $\forall$ are defined as abbreviations within this logic as usual. Furthermore we define the equality by $X = Y$ iff $X \subseteq Y \wedge Y \subseteq X$. A **sentence** is a formula without free variables. A formula resp. sentence of MSO is called a MSO-formula resp. MSO-sentence.

Note that 'monadic' refers to the fact that the only second-order quantification that is allowed is over monadic, i.e., unary predicates. Binary predicate symbols may occur in MSO-formulae, but only the unary ones may be quantified over.

For a given transition system $\mathcal{M}$ and an assignment $\beta : \mathrm{Var} \to \mathscr{P}\big(S^{\mathcal{M}}\big)$ the satisfaction relation $\models$ is defined inductively by:

$$
\begin{aligned}
(\mathcal{M}, \beta) &\models \mathrm{sr}(X) && \text{iff } \beta(X) = \{\mathrm{sr}^{\mathcal{M}}\},\\
(\mathcal{M}, \beta) &\models p(X) && \text{iff } \beta(X) \subseteq p^{\mathcal{M}},\\
(\mathcal{M}, \beta) &\models r(X, Y) && \text{iff there are } s, t \in S^{\mathcal{M}} \text{ such that } \beta(X) = \{s\},\\
& && \quad \beta(Y) = \{t\} \text{ and } (s, t) \in r^{\mathcal{M}},\\
(\mathcal{M}, \beta) &\models X \subseteq Y && \text{iff } \beta(X) \subseteq \beta(Y),\\
(\mathcal{M}, \beta) &\models \varphi \vee \psi && \text{iff } (\mathcal{M}, \beta) \models \varphi \text{ or } (\mathcal{M}, \beta) \models \psi,\\
(\mathcal{M}, \beta) &\models \neg\varphi && \text{iff not } (\mathcal{M}, \beta) \models \varphi,\\
(\mathcal{M}, \beta) &\models \exists X.\varphi(X) && \text{iff there is a } T \subseteq S^{\mathcal{M}} \text{ such that}\\
& && \quad (\mathcal{M}, \beta[X := T]) \models \varphi(X),
\end{aligned}
$$

where $\beta[X := T]$ denotes the assignment such that $\beta[X := T](X) = T$ and $\beta[X := T](Y) = \beta(Y)$ for $Y \neq X$.

For a MSO-sentence φ we write $\mathcal{M} \models \varphi$ if $(\mathcal{M}, \beta) \models \varphi$ is true for an arbitrary assignment. A MSO-sentence φ defines a class of transition systems by

$$\mathcal{C}^{\mathrm{MSO}}(\varphi) := \{\, \mathcal{M} \mid \mathcal{M} \text{ is a transition system and } \mathcal{M} \models \varphi \,\}.$$

Let $\mathcal{C}$ be a class of transition systems. $\mathcal{C}$ is **MSO-definable** if there is a MSO-sentence φ defining the class, i.e., $\mathcal{C} = \mathcal{C}^{\mathrm{MSO}}(\varphi)$ holds.

Remark 14.7. Not all MSO-definable classes of transition systems are bisimulation closed. Consider for example the MSO-sentence

$$\varphi := \exists X \exists Y.\big(\mathrm{sr}(X) \wedge r(X,Y) \wedge \forall Z.\big(r(X,Z) \Longrightarrow Y = Z\big)\big).$$

The sentence φ states that there is exactly one r-successor of the source. The class $\mathcal{C}^{\mathrm{MSO}}(\varphi)$ cannot be bisimulation closed because a bisimulation relation cannot fix any number of r-successors, i.e., if there is a r-successor in one transition system then there is one in all bisimilar systems, but there could be arbitrary many.

In the following we repeat the definition of the μ-calculus (cf. Chapter 10).

Definition 14.8. The signature of the **modal $\boldsymbol{\mu}$-calculus $\mathbf{L}_{\boldsymbol{\mu}}$** over transition systems contains only unary predicate symbols from Prop, binary predicate symbols from Rel and variables from Var. Formulae are defined inductively by the following grammar. Let $p \in \mathrm{Prop}$, $r \in \mathrm{Rel}$ and $X \in \mathrm{Var}$:

- X,
- p,
- $\neg\varphi$ for any formula φ,
- $\varphi \vee \psi$ for any formulae φ and ψ,
- $\langle r \rangle \varphi$ for any formula φ,
- $\mu X.\varphi(X)$ for any formula $\varphi(X)$ where X occurs only positively, i.e., under an even number of negations.

The dual of the modality $\langle r \rangle$ is denoted by $[r]$ and defined by $[r]\varphi := \neg\langle r \rangle\neg\varphi$. A formula resp. sentence of L_μ is called a L_μ-formula resp. L_μ-sentence. For a given transition system $\mathcal{M}$ and an assignment $\beta : \mathrm{Var} \to \mathscr{P}\big(S^{\mathcal{M}}\big)$ we define inductively the set $\|\varphi\|^{\mathcal{M}}_\beta$ in which the L_μ-formula φ is true:

$$\begin{aligned}
\|X\|^{\mathcal{M}}_\beta &:= \beta(X), \\
\|p\|^{\mathcal{M}}_\beta &:= p^{\mathcal{M}}, \\
\|\neg\varphi\|^{\mathcal{M}}_\beta &:= S^{\mathcal{M}} - \|\varphi\|^{\mathcal{M}}_\beta, \\
\|\varphi \vee \psi\|^{\mathcal{M}}_\beta &:= \|\varphi\|^{\mathcal{M}}_\beta \cup \|\psi\|^{\mathcal{M}}_\beta, \\
\|\langle r \rangle\varphi\|^{\mathcal{M}}_\beta &:= \big\{\, s \in S^{\mathcal{M}} \;\big|\; \mathrm{scc}^{\mathcal{M}}_r(s) \cap \|\varphi\|^{\mathcal{M}}_\beta \neq \emptyset \,\big\}, \\
\|\mu X.\varphi(X)\|^{\mathcal{M}}_\beta &:= \bigcap \big\{\, T \subseteq S^{\mathcal{M}} \;\big|\; \|\varphi(X)\|^{\mathcal{M}}_{\beta[X:=T]} \subseteq T \,\big\}.
\end{aligned}$$

Notice that we already used the Knaster-Tarski Theorem in the definition of $\|\mu X.\varphi(X)\|$ (cf. Theorem 20.4 in Chapter 20): due to the restriction that X may only occur positively in $\varphi(X)$, the operation $T \mapsto \|\varphi(X)\|^{\mathcal{M}}_{\beta[X:=T]}$ is monotone with respect to subset inclusion and the (existing) least fixed point of this map is exactly $\|\mu X.\varphi(X)\|^{\mathcal{M}}_{\beta}$. Monotone maps also have greatest fixed points. This is denoted by $\nu X.\varphi(X)$ and defined as $\neg\mu X.\neg\varphi[X := \neg X]$.

For a L_μ-sentence φ we write $(\mathcal{M}, s) \models \varphi$ if $s \in \|\varphi\|^{\mathcal{M}}_{\beta}$ holds for an arbitrary assignment. A L_μ-sentence φ defines a class $\mathcal{C}^{L_\mu}(\varphi)$ of transition systems by

$$\mathcal{C}^{L_\mu}(\varphi) := \{\, \mathcal{M} \mid \mathcal{M} \text{ is a transition system and } \left(\mathcal{M}, \mathrm{sr}^{\mathcal{M}}\right) \models \varphi \,\}.$$

Let $\mathcal{C}$ be a class of transition systems. $\mathcal{C}$ is **L_μ-definable** if there is a L_μ-sentence φ defining the class, i.e., $\mathcal{C} = \mathcal{C}^{L_\mu}(\varphi)$ holds.

As opposed to the situation of MSO we have:

Proposition 14.9. *Every L_μ-definable class is bisimulation closed.*

Proof. Let $\mathcal{M}$ and $\mathcal{N}$ be two transition systems and let R be a bisimulation relation between $\mathcal{M}$ and $\mathcal{N}$. It is easy to see that for any L_μ-sentence φ and for all $s \in S^{\mathcal{M}}$, $t \in S^{\mathcal{N}}$ with $(s, t) \in R$ the following holds:

$$s \in \|\varphi\|^{\mathcal{M}}_{\beta} \Longleftrightarrow t \in \|\varphi\|^{\mathcal{N}}_{\beta^*},$$

where β^* is an assignment for $\mathcal{N}$ derived from the assignment β for $\mathcal{M}$. Because we have $\left(\mathrm{sr}^{\mathcal{M}}, \mathrm{sr}^{\mathcal{N}}\right) \in R$ it follows:

$$\left(\mathcal{M}, \mathrm{sr}^{\mathcal{M}}\right) \models \varphi \wedge \mathcal{N} \sim \mathcal{M} \Longrightarrow \left(\mathcal{N}, \mathrm{sr}^{\mathcal{N}}\right) \models \varphi.$$

Hence $\mathcal{C}^{L_\mu}(\varphi)$ is bisimulation closed. □

Remark 14.10. We consider only definability by sentences. For MSO it makes no difference because the quantification is available. But in the case of the μ-calculus it is a proper restriction. To show this we define for an arbitrary L_μ-formula φ:

$$\mathcal{C}^{L_\mu}_*(\varphi) := \{\, \mathcal{M} \mid \mathcal{M} \text{ is a transition system and } \mathrm{sr}^{\mathcal{M}} \in \|\varphi\|^{\mathcal{M}}_{\beta} \text{ for all assignments } \beta \,\}.$$

A class $\mathcal{C}$ of transition systems is called **L_μ-formula-definable** if there is a L_μ-formula φ defining the class. There are L_μ-formula-definable classes which are not closed under bisimulation. Consider for example the following L_μ-formula where $r \in \mathrm{Rel}$:

$$\varphi := \neg(\langle r\rangle X \wedge \langle r\rangle\neg X).$$

Let $\mathcal{C} := \mathcal{C}^{L_\mu}_*(\varphi)$. For a transition system $\mathcal{M}$ and an arbitrary assignment β we have $\mathrm{sr}^{\mathcal{M}} \in \|\varphi\|^{\mathcal{M}}_{\beta}$ iff either the set $\beta(X)$ or the complement of $\beta(X)$ does not contain any r-successor of $\mathrm{sr}^{\mathcal{M}}$, i.e., for all $\mathcal{M}$ in $\mathcal{C}$ we have that either

$\mathrm{scc}_r^{\mathcal{M}}(\mathrm{sr}^{\mathcal{M}}) \cap \beta(X)$ is empty or $\mathrm{scc}_r^{\mathcal{M}}(\mathrm{sr}^{\mathcal{M}})$ is a subset of $\beta(X)$ for every assignment $\beta(X)$. In particular we obtain for the special case $\beta(X) := \{s\}$ with $s \in S^{\mathcal{M}}$ that there is at most one $s \in S^{\mathcal{M}}$ such that $s \in \mathrm{scc}_r^{\mathcal{M}}(\mathrm{sr}^{\mathcal{M}})$ for every $\mathcal{M} \in \mathcal{C}$. But there are transition systems without this property although they are bisimilar to $\mathcal{M}$, so $\mathcal{C}$ is not bisimulation closed (cf. Remark 14.7).

One direction of the expressive completeness result is the following:

Proposition 14.11. *Every* L_μ*-definable class is* MSO*-definable as well.*

Proof. For every L_μ-formula φ there is a MSO-formula $\varphi^*(X)$ where the variable X does not occur in φ and such that for every transition system $\mathcal{M}$ and every assignment β with $\beta(X) = \{s\}$ for some $s \in S^{\mathcal{M}}$:

$$(\mathcal{M}, \beta) \models \varphi^*(X) \iff s \in \|\varphi\|_\beta^{\mathcal{M}}.$$

For that we define recursively:

- For $\varphi = Y$ let $\varphi^*(X) := X \subseteq Y$,
- For $\varphi = p$ let $\varphi^*(X) := \exists Y.\big(p(Y) \wedge X \subseteq Y\big)$,
- For $\varphi = \neg\psi$ let $\varphi^*(X) := \neg\psi^*(X)$,
- For $\varphi = \psi \vee \chi$ let $\varphi^*(X) := \psi^*(X) \vee \chi^*(X)$,
- For $\varphi = \langle r\rangle\psi$ let $\varphi^*(X) := \exists Y.\big(r(X,Y) \wedge \psi^*(Y)\big)$ where Y does not occur in ψ,
- For $\varphi = \mu Y.\psi(Y)$ let $\varphi^*(X)$ be a pure second order version of the statement

  $$\forall Y.\big(\forall z(z \in \psi^*(Y) \longrightarrow z \in Y) \longrightarrow X \subseteq Y\big),$$

  where z is an additional first order variable.

It is easy to check that φ^* satisfies the property above. We obtain for any assignment β:

$$\begin{aligned}\big(\mathcal{M}, \mathrm{sr}^{\mathcal{M}}\big) \models \varphi &\iff \mathrm{sr}^{\mathcal{M}} \in \|\varphi\|_\beta^{\mathcal{M}} \\ &\iff \big(\mathcal{M}, \beta[X := \{\mathrm{sr}^{\mathcal{M}}\}]\big) \models \varphi^*(X) \\ &\iff (\mathcal{M}, \beta) \models \exists X.\big(\mathrm{sr}(X) \wedge \varphi^*(X)\big).\end{aligned}$$

For an arbitrary L_μ-sentence φ the formula $\tilde{\varphi} := \exists X.\big(\mathrm{sr}(X) \wedge \varphi^*(X)\big)$ is a MSO-sentence. Hence it follows $\mathcal{C}^{\mathrm{L}_\mu}(\varphi) = \mathcal{C}^{\mathrm{MSO}}(\tilde{\varphi})$. □

14.4 μ-Automata and μ-Games

Definition 14.12. Let $\mathcal{U} = \{p_1, \ldots, p_n\}$ be a finite set of propositional letters and $\mathrm{Sent}(\mathcal{U})$ a set of sentences of the first order logic (possibly with equality predicate) over the signature consisting of the unary predicates $\{p_1, \ldots, p_n\}$. A **marking** of a set T is a function $m : \mathcal{U} \to \mathscr{P}(T)$.

In the sequel we consider structures of the form $(T, \{\, m(p) \mid p \in \mathcal{U} \,\})$, i.e., a structure with carrier T where each predicate $p \in \mathcal{U}$ is interpreted as $m(p)$. If a sentence $\varphi \in \mathrm{Sent}(\mathcal{U})$ is satisfied in this structure we write as usual

$$(T, \{\, m(p) \mid p \in \mathcal{U} \,\}) \models \varphi.$$

In our situation we fix a transition system $\mathcal{M}$. Let Q be a finite set of states (distinct from the ones of $S^{\mathcal{M}}$) and let $\Sigma_R \subseteq \mathrm{Rel}$ be a finite set of letters for relations. Then we let $\mathcal{U} := \Sigma_R \times Q$. So each predicate $p \in \mathcal{U}$ is of the form $p = (r, q)$ for $r \in \Sigma_R$ and $q \in Q$. We fix a state $s \in S^{\mathcal{M}}$ and consider the set T of the r-successors of s for all $r \in \Sigma_R$.

Example 14.13. Let $Q := \{q_1, q_2\}$ and $\Sigma_R := \{r\}$. Let $s \in S^{\mathcal{M}}$ and let $m(r, q_1)$ and $m(r, q_2)$ be sets of r-successors of s. In the structure $\mathcal{N}$ with carrier $\mathrm{scc}_r^{\mathcal{M}}(s)$ the predicate p_1 will be interpreted as $m(r, q_1)$ and p_2 as $m(r, q_2)$. We consider the first order formula $\varphi := \exists x_1, x_2. \big(\bigwedge_{i=1,2} p_i(x_i) \land \forall y. \bigvee_{i=1,2} p_i(y) \big)$. Then we have:

$$\mathcal{N} \models \varphi \iff m(r, q_1) \neq \emptyset \land m(r, q_2) \neq \emptyset \land m(r, q_1) \cup m(r, q_2) = \mathrm{scc}_r^{\mathcal{M}}(s).$$

So φ is true in $\mathcal{N}$ iff $m(r, q_1)$ and $m(r, q_2)$ forms a partition of the set of r-successors of s into two non-empty (but not necessarily disjunct) sets.

We are now ready to define the notion of μ-automata:

Definition 14.14. Let Q be a finite set of states and $q_{\mathrm{I}} \in Q$ an initial state. Let $\Sigma_P \subseteq \mathrm{Prop}$ be a finite set of unary predicate symbols, $\Sigma_R \subseteq \mathrm{Rel}$ a finite set of binary predicate symbols and $\delta : Q \times \mathscr{P}\big(\Sigma_P\big) \to \mathrm{Sent}(\Sigma_R \times Q)$ a transition function. Finally let $\Omega : Q \to \omega$ be a parity function defining the acceptance condition. Then we call the tuple

$$(Q, \Sigma_P, \Sigma_R, q_{\mathrm{I}}, \delta, \Omega)$$

an **$\boldsymbol{\mu}$-automaton $\boldsymbol{\mathcal{A}}$**.

In fact this is the definition of an alternating parity automaton. Observe that the μ-automaton has two alphabets Σ_P and Σ_R, the first is for checking properties of states and the second is for checking the labels of taken transitions.

We will define the acceptance of arbitrary transition systems by the μ-automaton in terms of games. But we introduce first some abbreviations: for a given transition system $\mathcal{M}$ and a state $s \in S^{\mathcal{M}}$ let

$$L^{\mathcal{M}}(s) := \big\{\, p \in \Sigma_P \;\big|\; (\mathcal{M}, s) \models p \,\big\} = \big\{\, p \in \mathrm{Prop} \;\big|\; s \in p^{\mathcal{M}} \,\big\} \cap \Sigma_P$$

be the set of all propositional letters p in Σ_P such that s satisfies p in $\mathcal{M}$ and

$$\mathrm{SCC}^{\mathcal{M}}(s) := \bigcup_{r \in \Sigma_R} \mathrm{scc}_r^{\mathcal{M}}(s)$$

the set of all r-successors of s for $r \in \Sigma_R$.

Definition 14.15. Let $\mathcal{M}$ be a transition system and let $\mathcal{A}$ be a μ-automaton. We consider the following **μ-game $\mathcal{G}(\mathcal{M}, \mathcal{A})$**:

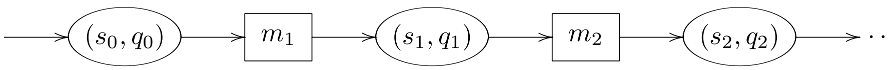

The initial position is $(s_0, q_0) = \big(\mathrm{sr}^{\mathcal{M}}, q_{\mathrm{I}}\big)$. If the current position is (s_i, q_i) then Player 0 is to move. Player 0 chooses a marking m_{i+1} of $\mathrm{SCC}^{\mathcal{M}}(s_i)$ – i.e., a function $m_{i+1} : \Sigma_R \times Q \to \mathscr{P}\big(\mathrm{SCC}^{\mathcal{M}}(s_i)\big)$ – such that:

- for every $r \in \Sigma_R$ and every $q \in Q$ the elements of $m_{i+1}(r, q)$ are r-successors of s_i,
- the structure $\mathcal{N} := \big(\mathrm{SCC}^{\mathcal{M}}(s_i), \{\, m_{i+1}(r, q) \mid r \in \Sigma_R, q \in Q \,\}\big)$ is a model of the first order sentence $\delta\big(q_i, L^{\mathcal{M}}(s_i)\big)$:

$$\mathcal{N} \models \delta\big(q_i, L^{\mathcal{M}}(s_i)\big).$$

If the current position is a marking m_i then Player 1 is to move and he chooses $r_i \in \Sigma_R$, $q_i \in Q$ and a state $s_i \in m_i(r_i, q_i)$. The pair (s_i, q_i) becomes the next position.

The criterion for winning the μ-game is as follows: one player wins if the other cannot make a move. Otherwise the play is infinite and we obtain the sequence

$$\big(\mathrm{sr}^{\mathcal{M}}, q_{\mathrm{I}}\big), m_1, (s_1, q_1), m_2, \dots$$

Let $\pi = q_{\mathrm{I}} q_1 q_2 \dots$ be the sequence of played states. Because of the finiteness of Q and the pigeonhole principle there is a $j \in \omega$ such that j appears infinitely often in the sequence $\Omega(q_{\mathrm{I}}), \Omega(q_1), \dots$ Let $\min \mathrm{Inf}(\Omega(\pi))$ be the smallest number with this property. Then Player 0 wins the μ-game iff $\min \mathrm{Inf}(\Omega(\pi))$ is even. So $\mathcal{G}(\mathcal{M}, \mathcal{A})$ is in fact a sort of a parity game.

The transition system $\mathcal{M}$ is **accepted by the μ-automaton $\mathcal{A}$** if there is a winning strategy f_0 for Player 0 in the μ-game $\mathcal{G}(\mathcal{M}, \mathcal{A})$. The class

$$L(\mathcal{A}) := \{\, \mathcal{M} \mid \mathcal{M} \text{ is a transition system accepted by } \mathcal{A} \,\}$$

is called the **language recognized by $\mathcal{A}$**.

We define for every $n \in \omega$ the formula **diff** of first order logic as follows:

$$\mathrm{diff}(x_1, \dots, x_n) := \bigwedge_{1 \le i < j \le n} x_i \neq x_j .$$

The formula "diff" states that the values of $x_1, \dots, x_n$ are pairwise different.

The main tool for our purpose is the following correspondence of μ-automata and formulae of L_μ and MSO respectively which was proved by Janin and Walukiewicz.

Theorem 14.16.

(1) *A class* $\mathcal{C}$ *of* **transition systems** *is* $\mathbf{L_\mu}$**-definable** *iff* $\mathcal{C} = L(\mathcal{A})$ *for a* μ*-automaton* $\mathcal{A} = (Q, \Sigma_P, \Sigma_R, q_I, \delta, \Omega)$ *such that* $\mathrm{Sent}(\Sigma_R \times Q)$ *contains only disjunctions of sentences of the form:*

$$\exists x_1, \ldots, x_m. \big(\bigwedge_{1 \le i \le m} p_{k_i}(x_i) \wedge \forall y. \bigvee_{1 \le i \le m} p_{k_i}(y) \big),$$

where $p_{k_i} \in \Sigma_R \times Q$ *for* $i \in \{1, \ldots, m\}$*;*

(2) *A class* $\mathcal{C}$ *of* **transition systems** *is* $\mathbf{L_\mu}$**-formula-definable** *iff* $\mathcal{C} = L(\mathcal{A})$ *for a* μ*-automaton* $\mathcal{A}\,(Q, \Sigma_P, \Sigma_R, q_I, \delta, \Omega)$ *such that* $\mathrm{Sent}(\Sigma_R \times Q)$ *contains only disjunctions of formulae of the form:*

$$\exists x_1, \ldots, x_m. \big(\bigwedge_{1 \le i \le m} p_{k_i}(x_i) \wedge \forall y. \chi(y) \big),$$

where $p_{k_i} \in \Sigma_R \times Q$ *for* $i \in \{1, \ldots, m\}$ *and* $\chi(y)$ *is a disjunction of conjunctions of formulae of the form* $p(y)$ *for* $p \in \Sigma_R \times Q$*;*

(3) *A class* $\mathcal{C}$ *of* **transition trees** *is* **MSO-definable** *iff* $\mathcal{C} = L(\mathcal{A})$ *for a* μ*-automaton* $\mathcal{A}\,(Q, \Sigma_P, \Sigma_R, q_I, \delta, \Omega)$ *such that* $\mathrm{Sent}(\Sigma_R \times Q)$ *contains only disjunctions of formulae of the form:*

$$\exists x_1, \ldots, x_m. \big(\bigwedge_{1 \le i \le m} p_{k_i}(x_i) \wedge \mathrm{diff}(x_1, \ldots, x_m) \; \wedge$$
$$\forall y. \big(\mathrm{diff}(y, x_1, \ldots, x_m) \longrightarrow \chi(y) \big) \big),$$

where $p_{k_i} \in \Sigma_R \times Q$ *for* $i \in \{1, \ldots, m\}$ *and* $\chi(y)$ *is a disjunction of conjunctions of formulae of the form* $p(y)$ *for* $p \in \Sigma_R \times Q$*.*

In all three cases empty disjunctions are allowed, i.e., the set $\mathrm{Sent}(\Sigma_R \times Q)$ may contain the sentence $\varphi = \bot$ (since we have $\bigvee \emptyset = \bot$).

In fact it can be shown that, if the μ-automaton $\mathcal{A}$ has the alphabets Σ_P and Σ_R then the corresponding formula which defines the class $\mathcal{C}$ is also in this language, i.e., only unary and binary predicate symbols of Σ_P and Σ_R respectively occur in the formula. The converse also holds: if $\mathcal{C}$ is defined by a formula φ such that the set of unary predicate symbols in φ is Σ_P and its set of binary predicate symbols is Σ_R then the corresponding μ-automata may be assumed to have the same alphabets Σ_P and Σ_R.

Item (1) is a reformulation of a result in [90] and the proof of item (2) can be found in [89]. For item (3) see Lemma 16.23 in Chapter 16.

Since most readers will not be familiar with μ-automata and the games played on them and since their transition function is unusual we will give an example here.

Example 14.17. We consider the L_μ-formula $\varphi := \langle r^* \rangle p$ which is equivalent to $\mu X.(p \vee \langle r \rangle X)$. So we have $(\mathcal{M}, s) \models \varphi$ iff there is a (possibly empty) r-path starting from s to a state in $\mathcal{M}$ where p holds. Let $\Sigma_P := \{p\}$, $\Sigma_R := \{r\}$ and $Q := \{q_1, q_2\}$. We define the μ-automaton

$$\mathcal{A} := (Q, \Sigma_P, \Sigma_R, q_1, \delta, \Omega),$$

where the parity function is defined as $\Omega(q_1) = 1, \Omega(q_2) = 0$ and the transition function as

$$\delta(q, P) := \begin{cases} \exists x_1, x_2.\big(\bigwedge_{i=1,2} p_i(x_i) \wedge \forall y. \bigvee_{i=1,2} p_i(y)\big) & \text{if } q = q_1 \text{ and } P = \emptyset, \\ \forall y.\bot \vee \exists x.\big(p_2(x) \wedge \forall y.p_2(y)\big) & \text{otherwise.} \end{cases}$$

Notice that since we have $\bigwedge \emptyset = \top$ and $\bigvee \emptyset = \bot$ the sentence

$$\exists x_1, \ldots, x_k.\big(\bigwedge_{1 \le i \le k} p_i(x_i) \wedge \forall y. \bigvee_{1 \le i \le k} p_i(y)\big) \tag{14.1}$$

is equivalent to $\forall y.\bot$ for $k = 0$. Hence the formulae $\delta(q, P)$ are disjunctions of sentences of the form (14.1) and therefore as stated in Theorem 14.16(1).

Let $\mathcal{M}$ be an arbitrary transition system. We consider the μ-game $\mathcal{G} := \mathcal{G}(\mathcal{M}, \mathcal{A})$. Notice that the game always starts with the position $(\mathrm{sr}^{\mathcal{M}}, q_1)$. If the current position in $\mathcal{G}$ is (s, q_j) for $j = 1, 2$ and $s \in S^{\mathcal{M}}$ then a move of Player 0 is a marking $m : \Sigma_R \times Q \to \mathscr{P}\big(\mathrm{scc}_r^{\mathcal{M}}(s)\big)$. This move is legal if the structure with carrier $\mathrm{scc}_r^{\mathcal{M}}(s)$ is a model of the formula $\delta(q_j, L^{\mathcal{M}}(s))$, where the predicate p_i is interpreted as the set $m(r, q_i)$ for $i = 1, 2$.

Claim. Assume that the current position in $\mathcal{G}$ is (s, q_2). Then Player 0 has a strategy to win.

Proof (of Claim). Player 0 plays the marking m defined as $m(r, q_1) = \emptyset$ and $m(r, q_2) = \mathrm{scc}_r^{\mathcal{M}}(s)$. We have to check that m is indeed a legal move and that this move leads Player 0 toward winning the game.
Case 1. There is no r-successor of s in $\mathcal{M}$. Then the structure with carrier $\emptyset$ is a model of $\forall y.\bot$, hence the move is legal. Since both $m(r, q_1)$ and $m(r, q_2)$ are empty Player 1 cannot respond with any position and looses the game.
Case 2. The set $\mathrm{scc}_r^{\mathcal{M}}(s)$ is non-empty. Since p_2 is interpreted as $m(r, q_2) = \mathrm{scc}_r^{\mathcal{M}}(s)$ we have

$$\big(\mathrm{scc}_r^{\mathcal{M}}(s), \{m(r, q_1), m(r, q_2)\}\big) \models \exists x.\big(p_2(x) \wedge \forall y.p_2(y)\big),$$

so the move is legal as well. Since $m(r, q_1)$ is empty Player 1 can only respond with a position (t, q_2) where t is a r-successor of s. To this position we can apply the same strategy again. If the resulting play is infinite then only q_2 is encountered infinitely often. So we have $\min \mathrm{Inf}(\Omega(\pi)) = 0$ and therefore Player 0 wins the game. $\square$(Claim)

Now we prove that $\mathcal{C}^{\mathrm{L}_\mu}(\varphi) = L(\mathcal{A})$.

($\subseteq$) Let $\mathcal{M}$ be a transition system with $\big(\mathcal{M}, \mathrm{sr}^{\mathcal{M}}\big) \models \varphi$, i.e., there is a sequence $s_0 = \mathrm{sr}^{\mathcal{M}}, s_1, \ldots, s_n$ with $s_{i+1} \in \mathrm{scc}_r^{\mathcal{M}}(s_i)$ for $i < n$ such that $(\mathcal{M}, s_n) \models p$. We may assume that $(\mathcal{M}, s_i) \not\models p$ for $i < n$.

If the current position is (s_i, q_1) with $i < n$ then Player 0 plays the marking m_{i+1} defined by $m_{i+1}(r, q_1) = \{s_{i+1}\}$ and $m_{i+1}(r, q_2) = \mathrm{scc}_r^{\mathcal{M}}(s_i)$. Since $L^{\mathcal{M}}(s_i) = \emptyset$ we have

$$\big(\mathrm{scc}_r^{\mathcal{M}}(s_i), \{m_{i+1}(r, q_1), m_{i+1}(r, q_2)\}\big) \models \delta(q_1, \emptyset),$$

so the move is legal. Then Player 1 must respond with the position (s_{i+1}, q_1), since otherwise he would loose the game by the claim above.

So eventually the position (s_n, q_1) is reached. Player 0 then plays the marking with $m_{n+1}(r, q_1) = \emptyset$ and $m_{n+1}(r, q_2) = \mathrm{scc}_r^{\mathcal{M}}(s_n)$. Now we have $L^{\mathcal{M}}(s_n) = \{p\}$ and the move is legal by

$$\big(\mathrm{scc}_r^{\mathcal{M}}(s_n), \{m_{n+1}(r, q_2), m_{n+1}(r, q_1)\}\big) \models \delta(q_1, \{p\}).$$

If Player 1 can make a move at all he can only respond with the position (t, q_2) for an r-successor t of s_n, so by the claim above he looses the game. This means that the strategy for Player 0 presented above is a winning strategy in the game $\mathcal{G}(\mathcal{M}, \mathcal{A})$ and therefore we obtain $\mathcal{M} \in L(\mathcal{A})$.

($\supseteq$) Let $\mathcal{M}$ be a transition system with $\big(\mathcal{M}, \mathrm{sr}^{\mathcal{M}}\big) \not\models \varphi$. Let (s, q_1) be the current position in the game $\mathcal{G}$. Since we have $\Sigma_R = \{r\}$ the states s_i of any prefix of a play in $\mathcal{G}$ form an r-path of $\mathcal{M}$ starting in $\mathrm{sr}^{\mathcal{M}}$. By the assumption we have $(\mathcal{M}, s) \not\models p$ and therefore $L^{\mathcal{M}}(s) = \emptyset$. Player 0 has to satisfy $\delta(q_1, \emptyset)$ in the structure with carrier $\mathrm{scc}_r^{\mathcal{M}}(s)$, so he must play two non-empty subsets $m(r, q_1)$ and $m(r, q_2)$ of $\mathrm{scc}_r^{\mathcal{M}}(s)$ such that the union is the whole set (cf. Example 14.13). Otherwise he would loose the game. If he can make a move at all then let $t \in m(r, q_1)$ be an r-successor of s. Player 1 responds with the position (t, q_1). By the assumption we have $(\mathcal{M}, t) \not\models p$ as well, so we can apply the same strategy again. With this strategy either Player 0 cannot make a move or an infinite game is played, where only q_1 is encountered infinitely often. Because $\Omega(q_1)$ is odd Player 1 wins the game. So we obtain a winning strategy for Player 1 and therefore we have $\mathcal{M} \notin L(\mathcal{A})$.

14.5 Expressive Completeness

Theorem 14.16 suggest a strong connection between monadic second order logic and the modal μ-calculus. But the basic sentences of MSO are more expressive. We are for example able to compare the number of r-successors of a state s with some constant by the use of the existential quantification together with the formula "diff$(\boldsymbol{x})$". On the other hand we conjecture that the equivalent μ-automaton for a MSO-definable and bisimulation closed class of transition systems should not use the formula "diff" and hence the class should be also L_μ-definable by the last theorem. In this section we will prove this conjecture. Notice that the considered μ-automata are non-deterministic, so the argument must deal with this fact, i.e., the μ-automaton may have only runs using instances of the formula "diff" but nevertheless the μ-automaton accepts a bisimulation closed class. This means that at last the acceptance of this class does not depend on the use of instances of the formula "diff" in the particular run.

Theorem 14.18. *Let $\mathcal{C}$ be a bisimulation closed class of transition systems, then*

$$\mathcal{C} \text{ is MSO-definable} \iff \mathcal{C} \text{ is } \mathrm{L}_\mu\text{-definable.}$$

For one direction we need the following lemma:

Lemma 14.19. *Let φ be a MSO-sentence. Then there is an effectively constructible L_μ-sentence $\widehat{\varphi}$ such that for every transition system $\mathcal{M}$:*

$$\widehat{\mathcal{M}} \models \varphi \iff \left(\mathcal{M}, \mathrm{sr}^{\mathcal{M}}\right) \models \widehat{\varphi}.$$

Before proving the lemma let us show how it implies the theorem:

Proof (of Theorem 14.18). Let $\mathcal{C}$ be a bisimulation closed class of transition systems.

($\Leftarrow$) By Proposition 14.11 every L_μ-definable class is MSO-definable as well.

($\Rightarrow$) Let φ be a MSO-sentence defining the class $\mathcal{C}$. Let $\mathcal{M}$ be an arbitrary transition system. By Exercise 14.2, $\mathcal{M}$ and $\widehat{\mathcal{M}}$ are bisimilar. Since $\mathcal{C}$ is bisimulation closed we obtain $\mathcal{M} \models \varphi \iff \widehat{\mathcal{M}} \models \varphi$. Let $\widehat{\varphi}$ be the L_μ-sentence given by Lemma 14.19, so

$$\mathcal{M} \models \varphi \iff \widehat{\mathcal{M}} \models \varphi \iff \left(\mathcal{M}, \mathrm{sr}^{\mathcal{M}}\right) \models \widehat{\varphi}.$$

In particular we have $\mathcal{C} = \mathcal{C}^{\mathrm{MSO}}(\varphi) = \mathcal{C}^{\mathrm{L}_\mu}(\widehat{\varphi})$ and so $\mathcal{C}$ is L_μ-definable. ☐

It remains to prove the lemma:

Proof (of Lemma 14.19). For a formula ψ of the form

$$\exists x_1, \ldots, x_m . \big(\bigwedge_{1 \le i \le m} p_{k_i}(x_i) \wedge \mathrm{diff}(x_1, \ldots, x_m) \wedge \\ \forall y . \big(\mathrm{diff}(y, x_1, \ldots, x_m) \longrightarrow \chi(y)\big)\big) \quad (14.2)$$

we define the formula ψ^* by substituting "true" for "diff" in ψ:

$$\psi^* := \exists x_1, \ldots, x_m . \big(\bigwedge_{1 \le i \le m} p_{k_i}(x_i) \wedge \forall y . \chi(y)\big). \quad (14.3)$$

For a disjunction $\theta = \psi_1 \vee \cdots \vee \psi_l$ let $\theta^* := \psi_1^* \vee \cdots \vee \psi_l^*$.

Let φ be a MSO-sentence and let $\mathcal{C}$ be the class of transition trees defined by φ (notice that we consider transition *trees* here). By Theorem 14.16(3) there is a μ-automaton $\mathcal{A} = (Q, \Sigma_P, \Sigma_R, q_{\mathrm{I}}, \delta, \Omega)$ such that $\mathcal{C} = L(\mathcal{A})$ and all formulae of $\mathrm{Sent}(\Sigma_R \times Q)$ have the form as stated in the theorem. In particular for every $q \in Q$ and $P \subseteq \Sigma_P$ the formula $\delta(q, P)$ is a disjunction of formulae of the form given by (14.2). Let $\delta^*(q, P) := (\delta(q, P))^*$. We define the μ-automaton $\mathcal{A}^*$ by

$$\mathcal{A}^* = (Q, \Sigma_P, \Sigma_R, q_{\mathrm{I}}, \delta^*, \Omega)\,.$$

Claim. Let $\mathcal{M}$ be a transition system. Then $\mathcal{M}$ is accepted by $\mathcal{A}^*$ iff $\widehat{\mathcal{M}}$ is accepted by $\mathcal{A}$.

Before proving the claim let us show how it implies the lemma. By definition of the function δ^* the μ-automaton $\mathcal{A}^*$ has the required form of Theorem 14.16(2). Hence there is a L_μ-sentence $\widehat{\varphi}$ such that $L(\mathcal{A}^*) = \mathcal{C}^{\mathrm{L}_\mu}(\widehat{\varphi})$. By Exercise 14.2 the ω-unravelling $\widehat{\mathcal{M}}$ is a transition tree for every transition system $\mathcal{M}$, hence we obtain by the claim

$$\widehat{\mathcal{M}} \in \mathcal{C} = L(\mathcal{A}) \Longleftrightarrow \mathcal{M} \in L(\mathcal{A}^*) = \mathcal{C}^{\mathrm{L}_\mu}(\widehat{\varphi}).$$

So we have

$$\widehat{\mathcal{M}} \models \varphi \Longleftrightarrow \bigl(\mathcal{M}, \mathrm{sr}^{\mathcal{M}}\bigr) \models \widehat{\varphi}.$$

It remains to prove the claim:

Proof (of Claim). ($\Rightarrow$) Suppose that $\mathcal{M}$ is accepted by $\mathcal{A}^*$. We want to show that $\widehat{\mathcal{M}}$ is accepted by $\mathcal{A}$. We consider the μ-games $\mathcal{G}^* := \mathcal{G}(\mathcal{M}, \mathcal{A}^*)$ and $\mathcal{G} := \mathcal{G}(\widehat{\mathcal{M}}, \mathcal{A})$. By the assumption Player 0 has a winning strategy f_0^* in the game $\mathcal{G}^*$. We want to define inductively a winning strategy f_0 for Player 0 in the game $\mathcal{G}$. For that we play the games $\mathcal{G}^*$ and $\mathcal{G}$ simultaneously and transfer each move of Player 1 from $\mathcal{G}$ to $\mathcal{G}^*$. Then we transfer the suggested move of Player 0 by the given strategy f_0^* in the game $\mathcal{G}^*$ back to $\mathcal{G}$. Both games have the initial position $\bigl(\mathrm{sr}^{\mathcal{M}}, q_{\mathrm{I}}\bigr)$. Let

$$\bigl(\mathrm{sr}^{\mathcal{M}}, q_{\mathrm{I}}\bigr), m_1, (u_1, q_1), \ldots, m_n, (u_n, q_n)$$

be a prefix of a play in the game $\mathcal{G}$ according to the induction. Since we have that u_{i+1} is an r_{i+1}-successor of u_i for some $r_{i+1} \in \Sigma_R$ we may assume that $u_{i+1} = u_i(a_{i+1}, r_{i+1}, s_{i+1})$ holds for every $i < n$ with $a_{i+1} \in \omega$ and $s_{i+1} \in S^{\mathcal{M}}$. Consider the corresponding prefix

$$\bigl(\mathrm{sr}^{\mathcal{M}}, q_{\mathrm{I}}\bigr), m_1^*, (s_1, q_1), \ldots, m_n^*, (s_n, q_n)$$

in the game $\mathcal{G}^*$ where s_{i+1} is an r_{i+1}-successor of s_i and the markings m_i^* are according to the strategy f_0^*. Let $m_{n+1}^* : \Sigma_R \times Q \to \mathscr{P}\bigl(\mathrm{SCC}^{\mathcal{M}}(s_n)\bigr)$ be the marking suggested by f_0^* for the current position (s_n, q_n). We define the marking $m_{n+1} : \Sigma_R \times Q \to \mathscr{P}\bigl(\mathrm{SCC}^{\widehat{\mathcal{M}}}(u_n)\bigr)$ by

$$m_{n+1}(r, q) := \bigcup_{a \in \omega} \bigl\{\, u_n(a, r, t) \bigm| t \in m_{n+1}^*(r, q) \,\bigr\}.$$

In particular we have

$$m_{n+1}^* \neq \emptyset \Longrightarrow m_{n+1} \neq \emptyset. \tag{14.4}$$

By definition of the ω-unravelling we have $m_{n+1}(r,q) \subseteq \mathrm{scc}_r^{\widehat{\mathcal{M}}}(u_n)$. Moreover it holds

$$s_n \in p^{\mathcal{M}} \Longleftrightarrow u_n \in p^{\widehat{\mathcal{M}}}, \tag{14.5}$$

in particular $L^{\mathcal{M}}(s_n) = L^{\widehat{\mathcal{M}}}(u_n)$ and therefore

$$\delta^*\big(q_n, L^{\mathcal{M}}(s_n)\big) = \big(\delta\big(q_n, L^{\widehat{\mathcal{M}}}(u_n)\big)\big)^*. \tag{14.6}$$

Next we define abbreviations for the two first order structures which occur in the rules of the games:

$$\mathcal{N} := \big(\mathrm{SCC}^{\mathcal{M}}(s_n), \{\, m^*_{n+1}(r,q) \mid r \in \Sigma_R, q \in Q \,\}\big)$$

and

$$\widehat{\mathcal{N}} := \Big(\mathrm{SCC}^{\widehat{\mathcal{M}}}(u_n), \{\, m_{n+1}(r,q) \mid r \in \Sigma_R, q \in Q \,\}\Big).$$

By the fact that m^*_{n+1} is a legal move of Player 0 in the game $\mathcal{G}^*$ we have

$$\mathcal{N} \models \delta^*\big(q_n, L^{\mathcal{M}}(s_n)\big). \tag{14.7}$$

Let ψ^* be some satisfied disjunct of $\delta^*\big(q_n, L^{\mathcal{M}}(s_n)\big)$ of the form (14.3). We will show that

$$\widehat{\mathcal{N}} \models \psi,$$

where ψ has the original form given by (14.2). By (14.4) the 'existential part' of ψ is satisfied by the structure $\widehat{\mathcal{N}}$ as well. Because of the ω-indexing there are infinitely many elements in $m_{n+1}(r,q)$ corresponding to each single element in $m^*_{n+1}(r,q)$. Hence we can always choose pairwise different witnesses in $\widehat{\mathcal{N}}$, i.e., the formula $\mathrm{diff}(x_1,\ldots,x_m)$ is additionally satisfied.

Next we check that $\widehat{\mathcal{N}}$ is a model of $\forall y.\chi(y)$ as well, in particular the restriction $\forall y.\big(\mathrm{diff}(y,x_1,\ldots,x_m) \longrightarrow \chi(y)\big)$ and therefore ψ is satisfied by $\widehat{\mathcal{N}}$. To see this let $v = u_n(a,r,t)$ be an arbitrary element of $\mathrm{SCC}^{\widehat{\mathcal{M}}}(u_n)$. Then t is an r-successor of s_n and by (14.7) we have $\mathcal{N} \models \chi(t)$, i.e., $\mathcal{N}$ is a model of some appropriate predicates $p(t)$ occurring in χ. Since each p is interpreted as $m^*_{n+1}(r,q)$ for some $q \in Q$ it follows that $t \in m^*_{n+1}(r,q)$ and therefore $v \in m_{n+1}(r,q)$ by the definition of m_{n+1}. This means that $\widehat{\mathcal{N}}$ is a model of the same predicates $p(v)$ and therefore a model of $\chi(v)$.

In summary this means that taking m_{n+1} is indeed a legal move of Player 0 in the game $\mathcal{G}$. So we define the value of the strategy f_0 for the current position by m_{n+1}. From this position Player 1 chooses some $r_{n+1} \in \Sigma_R$, $q_{n+1} \in Q$ and a state $u_{n+1} \in m_{n+1}(r_{n+1}, q_{n+1})$ with $u_{n+1} = u_n(a, r_{n+1}, t)$. The pair (u_{n+1}, q_{n+1}) becomes the next position in the game $\mathcal{G}$. Now we let $s_{n+1} := t$ and continue the game $\mathcal{G}^*$ by the move (s_{n+1}, q_{n+1}) of Player 1. We arrive at prefixes of plays in $\mathcal{G}$ and $\mathcal{G}^*$ satisfying our initial assumption.

It is clear that if Player 1 gets stuck in the game $\mathcal{G}^*$ then he cannot make a move in the game $\mathcal{G}$ as well. On the other hand by the inductive definition of the strategy f_0 Player 0 can always make a move in $\mathcal{G}$. Hence he cannot lose in a finite number of rounds. For an infinite play the result is the sequence

$$\big(\mathrm{sr}^{\mathcal{M}}, q_{\mathrm{I}}\big), m_1, (u_1, q_1), \ldots, m_n, (u_n, q_n), \ldots$$

The corresponding play in $\mathcal{G}^*$ is infinite as well:

$$\big(\mathrm{sr}^{\mathcal{M}}, q_{\mathrm{I}}\big), m_1^*, (s_1, q_1), \ldots, m_n^*, (s_n, q_n), \ldots$$

Let $\pi = q_{\mathrm{I}} q_1 q_2 \ldots$ be the sequence of the played automaton states, which is the same for both games. Because the play in $\mathcal{G}^*$ is according to the winning strategy f_0^* the smallest integer appearing infinitely often in the sequence $\Omega(q_{\mathrm{I}})\Omega(q_1)\ldots$ is even. But the parity function Ω is identical for both automata and therefore the value of $\min \mathrm{Inf}(\Omega(\pi))$ is the same. It follows that Player 0 wins the game $\mathcal{G}$ as well. Hence f_0 is indeed a winning strategy for Player 0 and $\widehat{\mathcal{M}}$ is accepted by the μ-automaton $\mathcal{A}$.

($\Leftarrow$) Suppose now that $\widehat{\mathcal{M}}$ is accepted by $\mathcal{A}$ and let f_0 be a winning strategy for Player 0 in the game $\mathcal{G}$. The argument is analogous to the one above with interchanged roles of the games, i.e., now we want to define inductively a winning strategy f_0^* for Player 0 in the game $\mathcal{G}^*$. We use the same notations as before. Let

$$\big(\mathrm{sr}^{\mathcal{M}}, q_{\mathrm{I}}\big), m_1^*, (s_1, q_1), \ldots, m_n^*, (s_n, q_n)$$

be a prefix of a play in the game $\mathcal{G}^*$ according to the induction and let

$$\big(\mathrm{sr}^{\mathcal{M}}, q_{\mathrm{I}}\big), m_1, (u_1, q_1), \ldots, m_n, (u_n, q_n)$$

be the corresponding prefix in the game $\mathcal{G}$ where we have: if $s_{i+1} \in \mathrm{scc}_r^{\mathcal{M}}(s_i)$ holds for $r \in \Sigma_R$ then $u_{i+1} = u_i(a, r, s_{i+1})$ for some $a \in \omega$. The markings m_i are played according to the strategy f_0.

Let $m_{n+1} : \Sigma_R \times Q \to \mathscr{P}\big(\mathrm{SCC}^{\widehat{\mathcal{M}}}(u_n)\big)$ be the marking suggested by f_0. We define the marking $m_{n+1}^* : \Sigma_R \times Q \to \mathscr{P}\big(\mathrm{SCC}^{\mathcal{M}}(s_n)\big)$ by

$$m_{n+1}^*(r, q) := \big\{\, t \in S^{\mathcal{M}} \;\big|\; \exists a \in \omega. \big(u_n(a, r, t) \in m_{n+1}(r, q)\big) \,\big\}.$$

Again we have $m_{n+1}^*(r, q) \subseteq \mathrm{scc}_r^{\mathcal{M}}(s_n)$ by the definition of the ω-unravelling. Since m_{n+1} is a legal move of Player 0 in the game $\mathcal{G}$ we have

$$\widehat{\mathcal{N}} \models \delta\big(q_n, L^{\widehat{\mathcal{M}}}(u_n)\big). \tag{14.8}$$

Let ψ be some satisfied disjunct of $\delta\big(q_n, L^{\widehat{\mathcal{M}}}(u_n)\big)$. We have to check that m_{n+1}^* is indeed a legal move of Player 0 in the game $\mathcal{G}^*$. By (14.6) it suffices to show that

$$\mathcal{N} \models \psi^*, \tag{14.9}$$

where ψ^* is the formula defined by (14.3). We may assume that the occurring predicates are $p_{k_i} = (r'_i, q'_i)$ with $r'_i \in \Sigma_R$ and $q'_i \in Q$ for every $i \in \{1, \ldots, m\}$. First we check that

$$m^*_{n+1}(r'_i, q'_i) \neq \emptyset \text{ for } i \in \{1, \ldots, m\}$$

and therefore

$$\mathcal{N} \models \exists x_1, \ldots, x_m. \bigwedge_{1 \leq i \leq m} p_{k_i}(x_i). \tag{14.10}$$

By (14.8) it follows that $\widehat{\mathcal{N}} \models \exists x_1, \ldots, x_m. \bigwedge_{1 \leq i \leq m} p_{k_i}(x_i)$, in particular $\widehat{\mathcal{N}} \models \exists x_i.p_{k_i}(x_i)$ for every $i \in \{1, \ldots, m\}$. Hence there is some $v \in \mathrm{SCC}^{\widehat{\mathcal{M}}}(u_n)$ such that $v \in m_{n+1}(r'_i, q'_i)$. By the definition of successors in $\widehat{\mathcal{M}}$ and the fact that $m_{n+1}(r'_i, q'_i)$ contains only r'_i-successors of u_n in $\widehat{\mathcal{M}}$ we have $v = u_n(a, r'_i, t)$ for some $a \in \omega$ and $t \in S^{\mathcal{M}}$. Hence $t \in m^*_{n+1}(r'_i, q'_i)$ by the definition of m^*_{n+1}. Next we check

$$\mathcal{N} \models \forall y.\chi(y). \tag{14.11}$$

Let $t \in \mathrm{scc}^{\mathcal{M}}_r(s_n)$ for some $r \in \Sigma_R$. We use again the property of the ω-unravelling that there are infinitely many different r-successors of u_n in $\widehat{\mathcal{M}}$ corresponding to each r-successor of s_n in $\mathcal{M}$. Hence there exists an $a \in \omega$ such that for $v = u_n(a, r, t)$ we have $\widehat{\mathcal{N}} \models \mathrm{diff}(v, x_1, \ldots, x_m)$. Therefore $\widehat{\mathcal{N}} \models \chi(v)$ holds by (14.8). Since χ is monotone in the predicates we obtain $\mathcal{N} \models \chi(t)$. To see this, notice that $\chi(v)$ has the form $\chi(v) = \bigvee_w \bigwedge_{w'} p_{w,w'}(v)$ with $p_{w,w'} \in \Sigma_R \times Q$. So we have $\widehat{\mathcal{N}} \models p_{w,w'}(v)$ for some appropriate pairs (w, w'), i.e., v is an element of $m_{n+1}(r, q_{w,w'})$. We obtain $t \in m^*_{n+1}(r, q_{w,w'})$ by the definition of m^*_{n+1} and therefore $\mathcal{N} \models p_{w,w'}(t)$ for the same predicates, i.e., $\mathcal{N} \models \chi(t)$ is true. By (14.10) and (14.11) and the definition of ψ^* we have $\mathcal{N} \models \psi^*$. This proves (14.9).

Taking m^*_{n+1} is therefore a legal move of Player 0 in the game $\mathcal{G}^*$. We define the value of the strategy f^*_0 for the current position by m^*_{n+1} and arrive at the prefix

$$\big(\mathrm{sr}^{\mathcal{M}}, q_{\mathrm{I}}\big), m_1, (u_1, q_1), \ldots, m_n, (u_n, q_n), m_{n+1}$$

in the game $\mathcal{G}$ and the corresponding prefix

$$\big(\mathrm{sr}^{\mathcal{M}}, q_{\mathrm{I}}\big), m^*_1, (s_1, q_1), \ldots, m^*_n, (s_n, q_n), m^*_{n+1}$$

in the game $\mathcal{G}^*$. From this position in $\mathcal{G}^*$ Player 1 chooses some $r_{n+1} \in \Sigma_R$, $q_{n+1} \in Q$ and a state $s_{n+1} \in m^*_{n+1}(r_{n+1}, q_{n+1})$ and the pair (s_{n+1}, q_{n+1}) becomes the next position in $\mathcal{G}^*$. By definition of m^*_{n+1} there is an $a \in \omega$ such that $u_n(a, r_{n+1}, s_{n+1}) \in m_{n+1}(r_{n+1}, q_{n+1})$. We choose some $a_{n+1} \in \omega$ with this property and define

$$u_{n+1} := u_n(a_{n+1}, r_{n+1}, s_{n+1}).$$

Then (u_{n+1}, q_{n+1}) is a legal move of Player 1 in the game $\mathcal{G}$ and we continue it by this move. Again we arrive at prefixes of plays in $\mathcal{G}$ and $\mathcal{G}^*$ satisfying our initial assumptions.

We have to check that f_0^* is indeed a winning strategy for Player 0 in the game $\mathcal{G}^*$. By the inductive definition of f_0^* Player 0 can always make a move and hence he cannot lose in a finite number of rounds. As in the first case the played automaton states in any infinite play of $\mathcal{G}$ and in the corresponding infinite play of $\mathcal{G}^*$ are the same and the parity functions of both automata are identical. Since the play in $\mathcal{G}$ is according to the winning strategy f_0 of Player 0 the value of $\min \mathrm{Inf}(\Omega(\pi))$ is even. It follows that Player 0 wins the game $\mathcal{G}^*$ as well. Hence f_0^* is indeed a winning strategy for Player 0 and $\mathcal{M}$ is accepted by the μ-automaton $\mathcal{A}^*$.

This completes the proof of the Claim and of Lemma 14.19. □

Since the branching time temporal logic CTL* is easily translatable into monadic second order logic over unwindings of transition systems and formulae resulting from this translation are bisimulation closed we obtain immediately a result of Dam shown in [44]:

Corollary 14.20. CTL* *is translatable into* L_μ.