Bisimulation-Invariant Logics: Beyond Finite (and Infinite)

— PhD Dissertation —

UNIVERSITY OF WARSAW
Faculty of Mathematics, Informatics and Mechanics

Jędrzej Kołodziejski

Supervisor
prof. Bartek Klin
University of Oxford

November 2023

Abstract

This dissertation is about logics that are invariant under bisimulation.

The first part investigates classical model-theoretic questions in the modal context. These include categoricity ("when does a set of sentences have a unique model?"), compactness ("does satisfiability of a set of sentences follow from satisfiability of its finite pieces?") and small model property ("how big model is needed to satisfy a set of sentences?"). The questions are asked for modal logic over several classes of models: all models, transitive models, two-way models and ordinal models. Simple characterizations expressed in terms of finiteness are given.

The second part introduces and investigates the countdown μ-calculus, an extension of the classical modal fixpoint logic with *countdown operators*. The new operators resemble the standard fixpoint operators. However, instead of referring to fixpoints they refer to the ordinal approximations of these fixpoints. The resulting logic expresses (un)boundedness properties such as "there exist arbitrarily long paths staring at a given point" that are inexpressible in the standard calculus. The classical correspondence with parity games and automata extends to suitably defined countdown games and automata. The connection is used to answer expressivity and decidability questions.

Streszczenie

Przedmiot niniejszej rozprawy stanowią logiki niezmiennicze na bisymulacje.

Część pierwsza poświęcona jest klasycznym problemom teoriomodelowym przeniesionym na grunt modalny. Do zagadnnień tych należą: kategoryczność ("kiedy dany zbiór zdań posiada unikatowy model?"), zwartość ("czy spełnialność zbioru zdań wynika ze spełnialności jego skończonych fragmentów?") i własność małego modelu ("jak duży model jest konieczny, by spełnić dany zbiór zdań?"). Pytania te stawiane są dla logiki modalnej nad różnymi klasami modeli: modelami dowolnymi, przechodnimi, dwukierunkowymi oraz porządkowymi. Udowodnione zostają proste charakteryzacje wyrażone w terminach skończoności.

Część druga wprowadza i bada odliczający rachunek μ – rozszerzenie klasycznego modalnego rachunku punktów stałych o *operatory odliczające*. Operatory te przypominają standardowe operatory stałopunktowe – zamiast jednak odnosić się do punktów stałych, odnoszą się do porządkowych aproksymacji tychże. Otrzymana w ten sposób logika wyraża własności (nie)ograniczoności niewyrażalne w standardowym rachunku takie jak "istnieją dowolnie długie ścieżki rozpoczynające się w danym punkcie". Wprowadzone zostają gry i automaty odliczające – uogólniające gry i automaty parzystości. Klasyczna odpowiedniość logiki, gier i automatów zostaje uogólniona do poziomu ich odliczających rozszerzeń. Uogólniona odpowiedniość jest następnie wykorzystana do badania siły wyrazu logiki oraz związanych z nią problemów decyzyjnych.

Acknowledgements

First and foremost I want to thank my supervisor, Bartek Klin. I am deeply grateful for the countless inspiring discussions we had and spot-on comments he gave me, both before and after he moved to Oxford. Without his patient help this dissertation would certainly not be possible, and his sense of elegance and smile made our cooperation not only fruitful but also pleasant. I would also like to thank Mikołaj Bojańczyk for all the brilliant and stimulating ideas he shared with me during the early stages of my PhD studies, which contributed greatly to their final outcome. I owe a lot to the community of the Automata Group in Warsaw, as well as logicians from the Department of Philosophy in Warsaw and from ILLC in Amsterdam. They shaped my scientific interests and taught me much. Special thanks go to Mikołaj Bojańczyk and Sławek Lasota for funding my research from their grants and the great freedom they gave me. Last but not least, I want to express my gratitude to my beloved family and dear friends for their great support, and to the University of Warsaw Choir for giving me so much needed breaks during my fascinating scientific journey.

Contents

Chapter 1

Introduction

Bisimulations. A notion central to this dissertation is that of a *bisimulation*. Bisimulations, as well as some related basic concepts, are formally introduced at the beginning of Subsection 2.1.1 of Chapter 2. Before we dive into the technical details, however, let us have a look at the motivating intuition and take a moment to appreciate the significance and beauty of the notion ([28] overviews the role of this fundamental concept in various fields and may serve as a more extensive introduction).

Analysis of various kinds of phenomena requires different tools and approaches. What is common is that before any investigations a decision must be made: what part of the information we take as relevant and what is ignored as accidental. If we are given a graph, the order in which its vertices appear is usually irrelevant. Similarly, we do not care about the specific names of functions in a program, the way a reasoning agent stores remembered facts, or the orientation of a depicted triangle. The reason for this is that the structure of all the listed entities does not depend on such details. Their various instances are *isomorphic*. Identification of isomorphic objects is uncontroversial to the point that it is often assumed without mentioning it explicitly.

Isomorphism, however, is an inherently *external* notion. Consider the following directed graphs with nodes colored either **red** or **blue**:

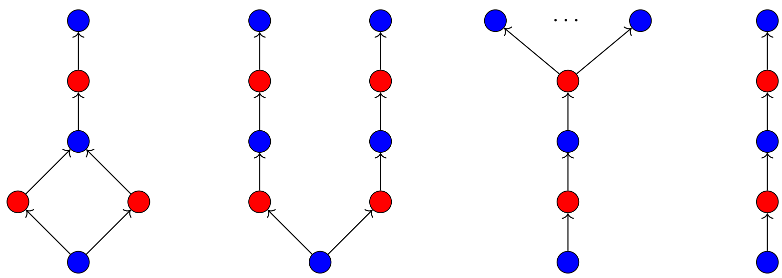

Any two of the above graphs are clearly non-isomorphic: the structure is different. However, to see the structural difference we need to look from the outside. On the other hand, in many contexts an *internal* perspective of an observer inhabiting the graphs is more appropriate than such an external one.

Imagine that the points of the graphs represent states of different programs controlling a spaceship. The **red** states are the ones in which the engine is on and the arrows depict possible changes of the state. Assume this is all we observe: we can only look at the engine or change the state to a next one. In such case every two states depicted in the same row exhibit the same *behavior*, in the sense that they cannot be distinguished by any present or future observations.

Similarly, assume that the points represent epistemic states of an agent, arrows depict their possible evolution and **blue** states are the ones in which (s)he believes that Socrates is mortal. Again, if in a point we can only see its color or move to a successor point then the observable behavior is identical for every two points depicted in the same row.

Such colored graphs, called *modal models*, describe an impressive number of seemingly unrelated phenomena. Notable examples other than programs and epistemic reasoning include: reasoning about possibility and necessity (hence the name *modal*), time and space, provability, but also non-well-founded set theory, or even moral duties (see [34] for an overview). What is common to all these different cases is that the properties of interest, such as "the engine can be always turned off", depend only on the behavior and not the full structure of a model.

A bisimulation captures such behavioral equivalence, similarly to how an isomorphism captures the structural one. As with an isomorphism, it is independent from the particular interpretation of modal models and allows us to abstract from it. The key idea here, formally introduced in Definition 2.1.4, resembles that of a congruence. Apart from having the same color, equivalent points m and m' must satisfy the so-called *back and forth* conditions: for every edge originating in either of the points there is an edge originating in the other one such that the targets of the two edges are again equivalent. The analogy to congruence is in fact deep and although we do not use category theory in this dissertation, it is worth to mention that in the category-theoretic framework bisimulations arise as a notion dual to homomorphisms between algebras [26].

Bisimulations successfully formalize the intuition about behavior and provide an elegant framework of both practical and theoretical significance. Although alternative formalizations are sometimes considered, bisimulation is the most popular and usually the most natural one (see [35] and [29] for a discussion). In all the listed examples of possible interpretations, the relevant properties of points in modal models are *invariant under bisimulation*: if two points are linked by a bisimulation then either both or none of them has a given property.

Logics. A natural approach to investigating various phenomena is to describe them using *logic*. The design of a good logic is often a nontrivial task because of the inevitable tradeoff: the more expressive power a logic has, the harder it is to reason about its formulae. Bisimulations provide an attractive solution, as they allow to filter out troublesome part of the information without loosing its essential aspects. This explains the great success of *bisimulation-invariant* logics, i.e. logics interpreted in points of modal models that only define properties invariant under bisimulation.

Among a great number of such logics two fundamental ones are *modal logic* ML and *modal μ-calculus* μ-ML. ML extends propositional logic with *modal operators* $\Diamond$ and $\Box$ working as restricted quantifiers. A formula $\Diamond\varphi$ interpreted in a point m means that m has a child satisfying φ. Dually, $\Box\varphi$ means that *all* m's children satisfy φ. The μ-calculus further extends ML with *fixpoints*.

Formulae of ML and μ-ML are constructed inductively in a way that guarantees bisimulation-invariance by design. However, they can be also obtained as semantically defined fragments of known logics. This is expressed by the two celebrated results: the van Benthem Theorem [33, Theorem 1.9] and the Janin-Walukiewicz Theorem [20, Theorem 11]. The first one identifies ML with the bisimulation-invariant fragment of the first-order logic FO, the fragment of FO consisting of precisely these formulae that define bisimulation-invariant properties. The latter theorem says that μ-ML is the same as the bisimulation-invariant fragment of monadic second-order logic MSO.

This demonstrates the usefulness of bisimulations. The full logics FO and MSO are rather complicated. For instance, the satisfiability problem (given a formula of the logic, does it have a model?) is undecidable for both. However, if we consider their bisimulation-invariant fragments ML and μ-ML then the question becomes decidable in both cases (as follows from [20]).

Games. Bisimulation-invariant logics ML and μ-ML and bisimulations themselves are easy to work with and theoretically pleasant. What contributes to that is the tight connection with *games*. The games we consider are always two-player zero-sum games, meaning that they are played between two players ∃ve and ∀dam, and precisely one of them wins. Games characterize the existence of bisimulations between models and the semantics of formulae. They offer a more dynamic point of view on both. Toggling between the the static, compositional perspective and the dynamic, game-theoretic one is often extremely helpful.

Automata. The connection between logic and games is mediated by *automata*. Given a model, an automaton induces a game and depending on its winner the model is either accepted or rejected. The interrelation with logic goes both ways. On the one hand, a formula φ can be viewed as an automaton $\mathcal{A}_\varphi$ equivalent to φ , meaning that $\mathcal{A}_\varphi$ accepts a point in a model iff φ is true there. On the other hand, an arbitrary abstract automaton $\mathcal{A}$ can be described by an equivalent formula $\varphi_\mathcal{A}$. Hence, logical formulae can be viewed as a special normal form of automata.

Games provide a useful perspective on logic but the converse is also true. The threefold connection between logic, games and automata is completed by *definability*. Classes of games corresponding to a given logic can be fully described in the logic itself, in the sense that there is always a formula that defines the class of all games in which ∃ve wins.

The contribution of this dissertation is naturally divided into two themes. Both extend known results about bisimulation-invariant logic in a way that can be roughly summed up as going beyond some form of *finiteness*. On top of that, we introduce a new notational framework for games, which contributes to the organization of present and possibly future proofs.

Model Theory for ML. First, we ask classical model-theoretic questions in the context of modal logic. Model theory for ML has been widely investigated to the point that it is virtually impossible to come up with a complete summary of all the results. The focus, however, was mostly put on model theory for *single formulae*. The theory for entire *sets* of formulae, although present, is less developed. While a finite set of formulae is equivalent to their conjunction and so the theory is the same, *infinite* sets behave quite differently. The *finite model property* could serve as an example. Every satisfiable ML formula has a *finite* model, but this is not true for arbitrary sets of formulae.

We scout this underexplored area in the landscape of modal logic. In particular, we investigate when a given set of modal formulae has a *unique* model. Such uniqueness property, known as *categoricity*, is usually understood up to isomorphism. However, in the context of ML bisimulation is arguably more natural than the isomorphism. This motivates the study of *bisimulational categoricity*, the property of having a unique model up to bisimulation. We investigate several variants of the question, when either all models, or only the ones from some fixed classes are considered. Simple characterizations phrased in terms of finite branching are found. We then discuss limitations of our method and turn to a fairly different class of models. For this last class, we establish characterizations of bisimulational categoricity and compactness as well as some results concerning the size of the models.

Countdown Logic. Second, we extend the modal fixpoint logic μ-ML. The classical μ-ML is already infinitary in some sense. It is capable of describing properties such as the existence (or lack) of infinite paths and the corresponding semantic games may contain infinite plays. However, μ-ML has its limitations. For instance, (un)boundedness properties such as "there exist arbitrarily long paths" are not expressible. This follows from the finite model property of μ-ML, because the conjunction "there are arbitrarily long paths *but* no infinite one" can be satisfied but only in infinite models.

Such (un)boundedness properties form a landmark in the quest for broadening our theoretic understanding. But apart from the more conceptual motivation, they are also important for interpretations. Whether every execution of a program terminates is often critical. However, the difference between a program with a bound on the maximal lengths of executions and one without such a bound is substantial, especially if one needs to allocate assets in advance.

A logic that was designed specifically to capture (un)boundedness is $\mathsf{MSO}+\mathsf{U}$ [8]. It extends MSO with a special nonstandard quantifier U interpreted as "there exist arbitrarily big finite sets such that...". Although complicated in general, MSO is well-behaved over some restricted classes of (finite or infinite) models such as words or binary trees. Over these classes MSO is equivalent to automata and thanks to this has decidable satisfiability, and the same was hoped for $\mathsf{MSO}+\mathsf{U}$ when it was introduced. Unfortunately, this is not the case. Unlike MSO, $\mathsf{MSO}+\mathsf{U}$ over trees or even infinite words has high *topological complexity* in a certain sense and consequently no simple game nor automata model may correspond to the logic [18, Theorem 5.1]. It eventually turned out that not only the classical proof techniques fail in the case of $\mathsf{MSO}+\mathsf{U}$, but its satisfiability problem is actually undecidable, even over infinite words [6, Theorem 1.1]. This discovery led to an end of the research project around $\mathsf{MSO}+\mathsf{U}$ in its original form, although some similar simpler logics (such as $\mathsf{WMSO}+\mathsf{U}$ in which second-order quantifiers range over *finite* sets) were successfully solved afterwards [3, Theorem 3 and 5], [7, Theorem 1].

Nonetheless, the question about a logic properly extending MSO but sharing its desirable properties remained open. Further research showed that, roughly, every nontrivial extension of the syntax of MSO leads to a logic containing $\mathsf{MSO}+\mathsf{U}$ and hence misses the goal [10, Theorem 1.3]. On the other hand, the good properties of MSO over words and trees stem from the somewhat miraculous equivalence between MSO and automata, which in turn correspond to μ-ML. This suggests a different approach: extend the syntax of μ-ML instead of that of MSO. We propose and investigate such an extension: the *countdown modal μ-calculus* $\mu^{<\infty}$-ML.

Recall that according to the Knaster-Tarski Theorem 2.1.1, the least and greatest fixpoint $\mathsf{LFP}.F$ and $\mathsf{GFP}.F$ of any monotone map $F : A \to A$ over a complete lattice A always exist. Moreover, they are given as the limits of the sequences of their inductively defined approximations:

$$f_\mu^\alpha = \bigvee_{\beta<\alpha} f(f_\mu^\beta)$$

and:

$$f_\nu^\alpha = \bigwedge_{\beta<\alpha} f(f_\nu^\beta)$$

with α ranging over ordinal numbers. The semantics of a logical formula φ can be seen as a map F on the powerset of the model. Under some assumptions on φ the map is monotone and so it has fixpoints $\mathsf{LFP}.F$ and $\mathsf{GFP}.F$. The classical μ-ML is obtained by enriching ML with operators μ and ν interpreted as these fixpoints. Our countdown calculus $\mu^{<\infty}$-ML extends μ-ML with additional *countdown* operators μ^α and ν^α for every ordinal α. Every such μ^α is similar to the classical μ except that it is interpreted as the α-th approximation F_μ^α of $\mathsf{LFP}.F$ instead of $\mathsf{LFP}.F$ itself. The semantics of ν^α is analogous.

This countdown logic $\mu^{<\infty}$-ML contains μ-ML by definition but is much more expressive. It can describe prototypical examples of (un)boundedness properties such as the existence of arbitrarily long paths in a model or arbitrarily long blocks of the same letter in an infinite word. Most importantly, however, we are able to define *countdown games* which lead to *countdown automata*. These extend the classical ones in a way such that the threefold logic-games-automata connection lifts to the new setting. Interestingly, the transgression of the finite-infinite dichotomy of μ-ML is accompanied by the loss of some of its finitary aspects. Countdown logic does not have the finite model property, countdown games are not memory-finite and countdown automata have infinitely many possible configurations. On the other hand, these are often *almost* finite in the sense that the internal symmetries allow for finite descriptions.

Some complications arise, however. First, the threefold correspondence requires a *vectorial* rather than a *scalar* calculus. Usually the two are equivalent thanks to the so-called *Bekić principle*. However, no analogous principle exists for countdown operators and the vectorial logic turns out to be strictly more expressive than its scalar fragment. Another difference is that, unlike in the standard setting, the automata model is inherently alternating and no nondeterministic model may exist. This is arguably a good sign, as the existence of a nondeterministic model would imply that the logic contains

$\mathsf{MSO} + \mathsf{U}$, and this is provably not the case for $\mu^{<\infty}$-ML due to its low topological complexity. But the lack of a nondeterministic model prevents us from directly adapting nice classical techniques. In particular, the model checking problem ("given a model $\mathcal{M}$ and a formula φ, does $\mathcal{M}$ satisfy φ?") is decidable, but satisfiability is solved only for some fragments and left as a conjecture for the full logic. Still, games and automata are crucial for establishing our decidability results as well as some model-theoretic properties of the logic.

Enhanced Bookkeeping. Apart from the more concrete results mentioned above, we introduce a new notational framework for games. It allows us to separate the coinductive heart of many intuitively similar proofs and replace repetitive vague claims with a reference to a formal statement. As such, it contributes to a better organization of the known and possibly also future proofs.

Organization of the Document. The dissertation is organized into five chapters. After this introduction 1, we set basic notions in Chapter 2. We then present the results related to model theory for ML in Chapter 3 and in Chapter 4 we introduce and investigate the countdown μ-calculus $\mu^{<\infty}$-ML. The last Chapter 5 contains a brief summary with a focus on possible further investigations. A significant part of the presented results was published by the author in [22] and [23]. Apart from that, the content of the document is new and was never published.

Chapter 2

Basic Notions

In this chapter we introduce basic notions used throughout the dissertation. This is mostly notation-setting, although the choices we made here facilitate further presentation and emphasize similarities between classical and new concepts. The only exception is Section 2.2 about games where we diverge from the usual terminology in two ways. Both are relevant mostly for Chapter 4, as games are nearly absent from the other Chapter 3.

First, we call the basic information unit of a game its *configuration* rather than a *position* (as it is usually called). Consequently, we talk about "configurational" rather than "positional" strategies etc. A configuration of a game may contain a position, but possibly also some extras such as counter values. The distinction helps us to better understand the connection between the usual parity games and countdown games introduced in Chapter 4.

Second, we introduce a novel notational framework for games. It allows us to avoid huge part of the usual hand-waving, both in the known and the new proofs. Since pedantic proofs are often good to have but painful to read, in order to keep the text reader-friendly we usually present the formal reasoning in parallel with its underlying intuition. Our new framework is based on the notion of a *partial game* where some moves lead to a *draw* meaning that the game ends but no player wins. Partial games generalize the games characterizing semantics of formulae, where the victory may depend on an external coloring of the model. A draw corresponds to reaching a configuration whose status is not fully determined in the sense that it depends on such an external coloring. However, the definition of a partial game is purely game-theoretic and abstracts from logic or automata.

2.1 Models

A *signature* is a set Symb whose elements, called *relational symbols* and *functional symbols*, come equipped with *arities* given by $\mathsf{ar} : \mathsf{Symb} \to \omega$. A *model* $\mathcal{M}$ for signature Symb consists of a *universe* M being a non-empty set, together with an *interpretation* $R^{\mathcal{M}} \subseteq M^{\mathsf{ar}(R)}$ for every relational symbol $R \in \mathsf{Symb}$ and $f^M : M^{\mathsf{ar}(f)} \to M$ for every functional symbol $f \in \mathsf{Symb}$. We will skip the superscript whenever the model is clear from the context, denote models by $\mathcal{M}$ and $\mathcal{N}$ and their elements, called *points*, by m and n. We will often abuse notation and identify model with its universe, writing $\mathsf{m} \in \mathcal{M}$ instead of $\mathsf{m} \in M$.

Binary Relations and Orderings. Assume a binary relation R. Whenever $\mathsf{m}R\mathsf{n}$, we call n an *R-successor* or *R-child* of m and m an *R-predecessor* or *R-parent* of n. We use standard terminology regarding properties of binary relations such as *transitive relation*, *reflexive relation* etc. The only possibly confusing property is *well-foundedness*: relation R is said to be *well-founded* if there is no infinite chain of elements $\mathsf{m}_1 R \mathsf{m}_2 R ...$ (with possible repetitions, so for example there is no m for which $\mathsf{m}R\mathsf{m}$). We call a single point m well-founded with respect to R (or just well-founded in case R is clear from the context) if there is no infinite chain originating in that m.

Throughout the thesis we will often use the notion of *partial orders*, i.e. binary relations $\preceq$ that are reflexive, antisymmetric and transitive. An ordering $\preceq$ is *linear* if any two elements m and m' are comparable (meaning that either $\mathsf{m} \preceq \mathsf{m}'$ or $\mathsf{m} \succeq \mathsf{m}'$) and *well-founded* if there is no infinite descending chain $\mathsf{m}_1 \succ \mathsf{m}_2 \succ ...$ (with $\mathsf{m} \prec \mathsf{m}'$ meaning $\mathsf{m} \preceq \mathsf{m}'$ *and* $\mathsf{m} \neq \mathsf{m}'$). It should be emphasized that well-foundedness of $\preceq$ *as an ordering* is the same as well-foundedness of the relation $\succ$ rather than $\preceq$ seen as a relation.

Lattices and Fixpoints. An important example of orderings are *ordinal numbers* or *ordinals*. The class of all ordinals will be denoted by Ord and the class of ordinals extended with a single additional element ∞, greater than all the ordinals, by Ord_∞. Recall that every well-founded linear ordering is isomorphic with an ordinal. We use the set-theoretic convention where the universe of every ordinal equals the set of all smaller ordinals. We use standard notation for ordinal arithmetic with $+1$, $+$ and $\times$ denoting successor, addition and multiplication, respectively.

A *complete lattice* is a partial ordering $(A, \preceq)$ such that for every $B \subseteq A$, the supremum and infimum of B, denoted $\bigvee B$ and $\bigwedge B$, respectively, both exist. Note that if we instantiate $B = A$, then $\bigvee A$ and $\bigwedge A$ are the greatest and the least elements of A, respectively. We denote these elements as $\top$ and $\bot$. The powerset $\mathcal{P}(X)$ of any set X ordered by set inclusion, or the unit interval $[0,1]$ with its usual ordering are examples of complete lattices. A function $f : A^k \to A$ is *monotone* if:

$$a_1 \preceq b_1 \wedge ... \wedge a_k \preceq b_k \text{ implies } f(a_1, ...a_k) \preceq f(b_1, ..., b_k).$$

For instance, monotonicity of an operation $f : \mathcal{P}(X) \to \mathcal{P}(X)$ means that $S \subseteq S'$ implies $f(S) \subseteq f(S')$. A classical result due to Knaster and Tarski [32] states that monotone operations in a complete lattice have both the least and the greatest fixpoint.

Theorem 2.1.1 (Knaster-Tarski)**.** *Assume that $f : A \to A$ is a monotone operation on a complete lattice $(A, \preceq)$. Then f has the greatest and the least fixpoint, denoted $\mathsf{GFP}.f$ and $\mathsf{LFP}.f$, respectively. Moreover,*

- $\mathsf{LFP}.f$ *is the limit of the increasing sequence:*
$$f_\mu^\alpha = \bigvee_{\beta<\alpha} f(f_\mu^\beta)$$
- $\mathsf{GFP}.f$ *is the limit of the decreasing sequence:*
$$f_\nu^\alpha = \bigwedge_{\beta<\alpha} f(f_\nu^\beta)$$

where $\alpha \in \mathsf{Ord}$. Hence, we denote $f_\mu^\infty = \mathsf{LFP}.f$ and $f_\nu^\infty = \mathsf{GFP}.f$.

Well-foundedness and Depth. An important instance of a fixpoint is the set M_{WF} of all well-founded points in a fixed structure (M, R). Consider the monotone operation $F_\Box : \mathcal{P}(M) \to \mathcal{P}(M)$:

$$F_\Box(H) = \{\mathsf{m} \in M \mid \forall_{\mathsf{m}R\mathsf{n}}\ \mathsf{n} \in H\}$$

mapping each H to the set of points whose all R-children belong to H. The least fixpoint of $F_\Box$ equals to the set of well-founded points:

$$\mathsf{LFP}.F_\Box = M_{\mathsf{WF}}. \tag{2.1}$$

We skip the easy proof and only illustrate the equality with an example.

Example 2.1.2. Consider the following model (M, R) with a single binary relation R depicted by arrows and **red** nodes representing the subsets ${F_\Box}_\mu^0 \subseteq {F_\Box}_\mu^1 \subseteq ... \subseteq {F_\Box}_\mu^\alpha \subseteq ... \subseteq M$:

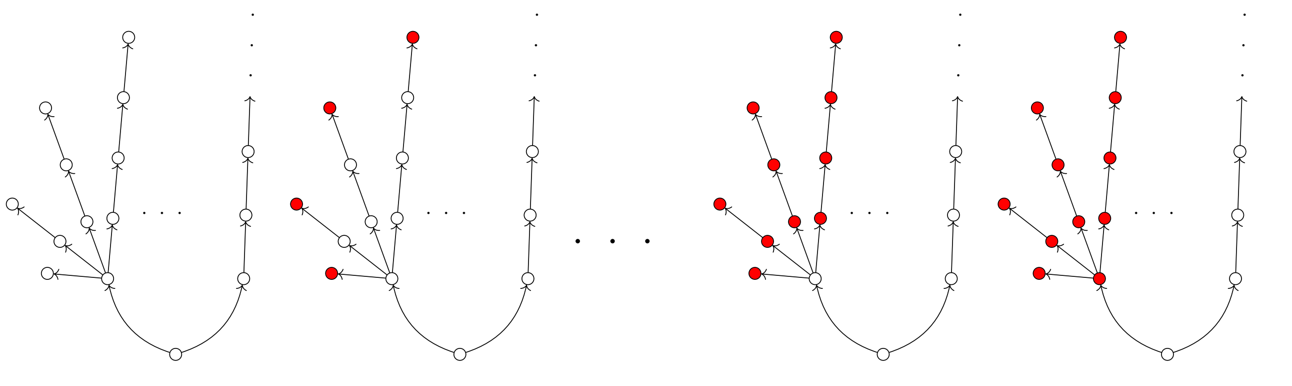

${F_\Box}_\mu^0 = \emptyset$ ${F_\Box}_\mu^1 = F_\Box(\emptyset)$ ${F_\Box}_\mu^\omega = \bigcup_{\alpha<\omega} {F_\Box}_\mu^\alpha$ ${F_\Box}_\mu^{\omega+1} = F_\Box({F_\Box}_\mu^\omega)$

The computation starts with ${F_\Box}^0_\mu = \bigvee\emptyset = \emptyset$. For every $\alpha < \omega$, the set ${F_\Box}^{\alpha+1}_\mu = F_\Box({F_\Box}^\alpha_\mu)$ contains points from which there is no path longer than $\alpha + 1$ (with the origin of a path contributing to its length). In the first limit step we get the set $F^\omega_\mu = \bigcup_{\alpha<\omega} {F_\Box}^\alpha_\mu$ of all points for which there exist a finite bound on the lengths of paths originating there. Finally, the set ${F_\Box}^{\omega+1}_\mu = F_\Box({F_\Box}^\omega_\mu) = \mathsf{LFP}.F_\Box$ containing all the well-founded points is the fixpoint because $F_\Box({F_\Box}^{\omega+1}_\mu) = {F_\Box}^{\omega+1}_\mu$. A different model could require a different number of steps to reach $\mathsf{LFP}.F_\Box$ but it always coincides with the set of well-founded points.

The fixpoint characterization gives rise to an ordinal-valued measure of well-foundedness expressed by the partial function:

$$\mathsf{depth} : M \to \mathsf{Ord}$$

mapping each well-founded m to the least α such that $\mathsf{m} \in {F_\Box}^\alpha_\mu$. The $\mathsf{depth}(\mathsf{m})$ need not be finite, but when it is it equals to the maximal length of paths starting in m.

For example, the nesting of operators from a fixed set O in a logical formula φ can be defined as the maximal depth of points in (M, R), where M is the set of all subformulae of φ beginning with an operator from O and $\psi R \psi'$ iff ψ' is a subformula of ψ (the nesting is 0 if there are no operations from O and 1 if they appear but are not nested).

2.1.1 Modal Models

We call a signature *modal* if it consists only of unary relational symbols Prop and binary relational symbols $\{\xrightarrow{\mathsf{a}} \mid \mathsf{a} \in \mathsf{Act}\}$, where Prop and Act are two sets called *atomic propositions* and *actions*, respectively. We call the relations $\xrightarrow{\mathsf{a}}$ *accessibility relations* and refer to $\xrightarrow{\mathsf{a}}$ as an a-*edge*. A *Kripke model* is a model for a modal signature. A *color* of a point in a Kripke model is the set of all atomic propositions it satisfies. In depictions of modal models we will often use actual colors (e.g. **blue** or **red**) to represent colors of points. Signatures and models are *monomodal* if there is only one action $\mathsf{a} \in \mathsf{Act}$, in which case we skip the index and write $\to$ in place of $\xrightarrow{\mathsf{a}}$. In either case, these symbols should not be confused with implication which is always denoted by $\implies$. Hence, "$\mathsf{m} \to \mathsf{n} \implies \mathsf{n} \to \mathsf{m}$" is a statement about monomodal $\to$: "if there is an edge from m to n then there is an opposite one from n to m".

A *pointed* model is a Kripke model $\mathcal{M}$ together with a chosen *root* $\mathsf{m} \in M$. Following the traditions of modal logic, we denote such pointed model by $\mathcal{M}, \mathsf{m}$ with no parentheses. The main focus of this thesis are logics and phenomena that are, in a broad sense, modal. Therefore, we will always assume that models and signatures are modal, unless it is explicitly stated otherwise.

Remark 2.1.3. The models we consider are both *edge-* and *vertex-labelled.* Usually, these two ways of coloring are equivalent in the sense that model-theoretic and computational questions about models with colors exclusively on edges or exclusively on vertices (and sometimes with no colors at all) are inter-reducible via straightforward coding. Although this is often true in our investigations, in some cases the situation is more subtle (see e.g. Conjecture 4.9.2). Moreover, some notions are more naturally presented in a mixed setting (as in Chapter 3, where it is arguably natural to see the colors of edges as fixed and the ones on vertices as a parameter). Therefore, we choose such richer, although sometimes redundant, notion of a Kripke model.

It is worth to mention that this notion can be further extended to the category-theoretic concept of a *coalgebra.* Coalgebras neatly generalize various types of structures such as Kripke models, (in)finite words, weighted graphs and many others. We do not assume familiarity with category theory and hence refrain ourselves from using coalgebras (an excellent introduction can be found in [26] or [19]). However, a huge part of what we present generalizes to that framework. We will comment on that in the concluding Chapter 5.

A key notion in this thesis is that of a *bisimulation.*

Definition 2.1.4. A a relation $Z \subseteq M \times M'$ between two (not necessarily distinct) models $\mathcal{M}$ and $\mathcal{M}'$ is called a *bisimulation* if for every $\mathsf{m}Z\mathsf{m}'$, the following three conditions are satisfied:

- (base) m and m' have the same color
  (that is: $\mathsf{m} \in \tau^{\mathcal{M}} \iff \mathsf{m}' \in \tau^{\mathcal{M}'}$ for every $\tau \in \mathsf{Prop}$);
- (forth) for every $\mathsf{a} \in \mathsf{Act}$: whenever $\mathsf{m} \overset{\mathsf{a}}{\to}^{\mathcal{M}} \mathsf{n}$, there exists n' such that $\mathsf{m}' \overset{\mathsf{a}}{\to}^{\mathcal{M}'} \mathsf{n}'$ and $\mathsf{n}Z\mathsf{n}'$;
- (back) for every $\mathsf{a} \in \mathsf{Act}$: whenever $\mathsf{m}' \overset{\mathsf{a}}{\to}^{\mathcal{M}'} \mathsf{n}'$, there exists n such that $\mathsf{m} \overset{\mathsf{a}}{\to}^{\mathcal{M}} \mathsf{n}$ and $\mathsf{n}Z\mathsf{n}'$.

Z is a bisimulation between pointed models $\mathcal{M}, \mathsf{m}$ and $\mathcal{M}', \mathsf{m}'$ (denoted $(\mathcal{M}, \mathsf{m})Z(\mathcal{M}', \mathsf{m}')$) if it is a bisimulation between $\mathcal{M}$ and $\mathcal{M}'$ such that $\mathsf{m}Z\mathsf{m}'$.

Bisimulations are closed under unions, and so the greatest bisimulation relation between given models, called *bisimilarity* and denoted $\leftrightarroweq$, always exists.

Example 2.1.5. Recall the monomodal example from the introduction. Roots, all depicted at the bottom, are indicated by arrows with no sources:

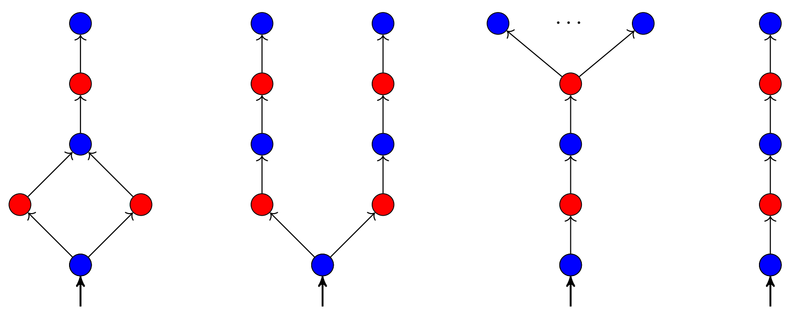

Any two of the above models are bisimilar. This is witnessed by the bisimulation relation Z that links two points iff they are at the same level.

Bisimilarity ignores a lot of information yet it preserves many important properties. For example, if points m and m' are bisimilar then m is well-founded iff so is m'. Moreover, in the case both are well-founded they have the same depth $\mathsf{depth}(\mathsf{m}) = \mathsf{depth}(\mathsf{m}')$. Let us recall three classical constructions that transform models into different but bisimilar ones.

Proposition 2.1.6 (generated submodels)**.** *Given a pointed model $\mathcal{M}, \mathsf{m}$, the model generated by m, denoted $\mathcal{M}_{\langle \mathsf{m} \rangle}$, is the submodel of $\mathcal{M}$ consisting of points reachable from m by a finite path. Then, the graph of the inclusion map $\iota : M_{\langle \mathsf{m} \rangle} \to M$ is a bisimulation.*

Proposition 2.1.7 (tree unravellings)**.** *For a pointed model $\mathcal{M}, \mathsf{m}$, its* tree unravelling *$\mathcal{M}^T, \mathsf{m}$ is the tree of all paths in $\mathcal{M}$ starting at m. That is, M^T is the set of all the paths $\mathsf{m}_1\mathsf{a}_1\mathsf{m}_2\mathsf{a}_2 ... \mathsf{m}_k \in (M \cup \mathsf{Act})^+$ such that $\mathsf{m}_1 = \mathsf{m}$ and $\mathsf{m}_i \overset{\mathsf{a}_i}{\to}^{\mathcal{M}} \mathsf{m}_{i+1}$ for every $i < k$. We interpret atomic propositions and accessibility relations as follows. For every $\pi, \pi' \in M^T$:*

- *for all $\mathsf{a} \in \mathsf{Act}$:*
$$\pi \overset{\mathsf{a}}{\to}^{\mathcal{M}^T} \pi' \quad \textit{iff} \quad \pi' = \pi \mathsf{a} \mathsf{n}$$
with $\mathsf{n} \in M$;
- *for all $\tau \in \mathsf{Prop}$:*
$$\pi \in \tau^{\mathcal{M}^T} \quad \textit{iff} \quad \pi \textit{ ends with } \mathsf{n} \textit{ such that } \mathsf{n} \in \tau^{\mathcal{M}}.$$

It follows that the graph of the function $\mathsf{last} : M^T \to M$ mapping each path $\pi \in M^T$ to its last point is a bisimulation.

Proposition 2.1.8 (quotients)**.** *Assume a model $\mathcal{M}$ and a bisimulation $Z \subseteq M \times M$ which happens to be an equivalence relation. Then, the model structure on the set of all equivalence classes of Z:*

- *for all* $\mathsf{a} \in \mathsf{Act}$*:*

$$[\mathsf{m}]_{/Z} \xrightarrow{\mathsf{a}}^{\mathcal{M}/Z} [\mathsf{n}]_{/Z} \quad \textit{iff} \quad \mathsf{m} \xrightarrow{\mathsf{a}}^{\mathcal{M}} \mathsf{n};$$

- *for all* $\tau \in \mathsf{Prop}$*:*

$$[\mathsf{m}]_{/Z} \in \tau^{\mathcal{M}/Z} \quad \textit{iff} \quad \mathsf{m} \in \tau^{\mathcal{M}}$$

is well-defined (i.e. does not depend on the choices of m *and* n*). Moreover, the graph of the projection map* $\mathsf{m} \overset{\pi_Z}{\longmapsto} [\mathsf{m}]_{/Z}$ *is a bisimulation. We call that model the* quotient *of* $\mathcal{M}$ *by* Z *and denote it* $\mathcal{M}_{/Z}$*.*

2.2 Games

Throughout the thesis we will often consider games, always meaning perfect-information games played between two players $\exists$ve and $\forall$dam (also denoted $\exists$ and $\forall$). Such a game $\mathcal{G}$ consists of a set Conf of configurations and a relation $\mathsf{Mov} \subseteq \mathsf{Conf} \times \mathsf{Conf}$ together with partitions $\mathsf{Conf}_\exists \sqcup \mathsf{Conf}_\forall = \mathsf{Conf}$ and $\mathsf{Win}_\exists \sqcup \mathsf{Win}_\forall = \mathsf{Conf}^\omega$ of configurations and their infinite sequences, respectively. A *play* is a (finite or infinite) sequence of configurations $\gamma_1\gamma_2...$ such that each consecutive γ_i and γ_{i+1} are linked by Mov. After a finite play $\pi = \gamma_1...\gamma_k$, the player P for which $\gamma_k \in \mathsf{Conf}_P$, called the *owner* of γ_k, has to extend the play by choosing γ_{k+1} with $\gamma_k \mathsf{Mov} \gamma_{k+1}$. Either at some point the game reaches a configuration where one of the players is stuck, meaning that (s)he has no legal choice, or it continues forever. In the first case the player who got stuck looses, and in the later P wins iff the resulting infinite play $\gamma_1\gamma_2\gamma_3...$ belongs to Win_P (thus, we call the partition $\mathsf{Win}_\exists \sqcup \mathsf{Win}_\forall$ the *winning condition* of the game). We say that the game moves *deterministically* from $\gamma \in \mathsf{Conf}$ if there is precisely one legal move from γ (in which case the ownership of γ does not matter). Usually, we will assume that the games are *initialized* in a fixed configuration γ. We denote the game $\mathcal{G}$ initialized at γ by $\mathcal{G}, \gamma$.

A *strategy* for player P is a function $\sigma : \mathsf{Conf}^*\mathsf{Conf}_P \to \mathsf{Conf}$ that tells the player how to play. A play π is *consistent* with strategy σ if whenever $\gamma_1...\gamma_k\gamma_{k+1}$ is a prefix of π and $\gamma_k \in \mathsf{Conf}_P$, then $\gamma_{k+1} = \sigma(\gamma_1...\gamma_k)$. We call such plays σ-plays and say that σ is *winning* from configuration γ if every σ-play starting at γ is won by P. Player *wins* the initialized game $\mathcal{G}, \gamma$ if (s)he has a strategy winning from γ.

A *phase* of a game is a set of its finite plays that is convex with respect to the prefix ordering, meaning that if π_1 is a prefix of π_2, π_2 is a prefix of π_3 and both π_1 and π_3 belong to the set then so does π_2. Given a phase $\mathcal{B}$ and a play $\pi \in \mathcal{B}$, we denote by $\mathcal{B}_\pi$ the subset of $\mathcal{B}$ consisting of all the plays having π as a prefix.

Bisimilarity Game. An important example of a game is the *bisimilarity game* $\mathcal{G}_{\leftrightarroweq}(\mathcal{M},\mathcal{M}')$ for models $\mathcal{M}$ and $\mathcal{M}'$, defined as follows. Players move two pebbles placed in points of $\mathcal{M}$ and $\mathcal{M}'$, round by round. At the beginning of each round, we check if the points where the pebbles are placed satisfy the base condition from Definition 2.1.4, i.e. they have the same color. If this is not the case then ∃ve looses immediately. Otherwise, ∀dam checks either the back or the forth condition: first he chooses an action $\mathsf{a} \in \mathsf{Act}$ and moves one pebble along a chosen a-edge and then ∃ve has to respond by moving the other pebble along an a-edge. In case of an infinite play, ∃ve wins.

Formally, we define the set of configurations to be:

$$\begin{aligned}\mathsf{Conf}_\exists =& M \times M' \times \{\mathsf{base}\} \sqcup M \times M' \times \mathsf{Act} \times \{\mathsf{back}, \mathsf{forth}\}\\ \mathsf{Conf}_\forall =& M \times M' \times \{\mathsf{b\&f}\}.\end{aligned}$$

The legal moves are as follows:

1. In $(\mathsf{m}, \mathsf{m}', \mathsf{base})$, if m and m' have different colors then there are no moves (so ∃ve looses immediately), otherwise the game moves deterministically to $(\mathsf{m}, \mathsf{m}', \mathsf{b\&f})$.
2. In $(\mathsf{m}, \mathsf{m}', \mathsf{b\&f})$, ∀dam chooses either:
   - $(\mathsf{n}, \mathsf{m}', \mathsf{a}, \mathsf{forth})$ such that $\mathsf{m} \xrightarrow{\mathsf{a}}^{\mathcal{M}} \mathsf{n}$ or
   - $(\mathsf{m}, \mathsf{n}', \mathsf{a}, \mathsf{back})$ such that $\mathsf{m}' \xrightarrow{\mathsf{a}}^{\mathcal{M}'} \mathsf{n}'$.
3. In $(\mathsf{n}, \mathsf{m}', \mathsf{a}, \mathsf{forth})$ or $(\mathsf{m}, \mathsf{n}', \mathsf{a}, \mathsf{back})$, ∃ve chooses $(\mathsf{n}, \mathsf{n}', \mathsf{base})$ such that $\mathsf{m}' \xrightarrow{\mathsf{a}}^{\mathcal{M}'} \mathsf{n}'$ or $\mathsf{m} \xrightarrow{\mathsf{a}}^{\mathcal{M}} \mathsf{n}$, respectively.

If during the play one of the players is stuck, (s)he looses immediately and otherwise ∃ve wins. It is well known that the above game characterizes bisimilarity, meaning that for every $\mathsf{m} \in M$ and $\mathsf{m}' \in M'$:

$$\exists\text{ve wins } \mathcal{G}_{\leftrightarroweq}(\mathcal{M},\mathcal{M}') \text{ from } (\mathsf{m}, \mathsf{m}', \mathsf{base}) \iff \mathcal{M}, \mathsf{m} \leftrightarroweq \mathcal{M}', \mathsf{m}'. \tag{2.2}$$

A standard proof can be found e.g. in [37].Later in this chapter, in Example 2.2.10, we also derive (2.2) from a more general result.

Depth-k Bisimilarity Game. Another important example of a game, this time with configurations containing a *counter value*, is the game characterizing the relation $\leftrightarroweq^k$ of *depth-k bisimilarity.* This relation, defined for every $k < \omega$, approximates bisimilarity and captures the intuition that models are *bisimilar up to depth k*. Instead of defining depth-k bisimulations (which is also possible), we directly define relations of $\leftrightarroweq^0 \supseteq \leftrightarroweq^1 \supseteq \ldots \supseteq \leftrightarroweq^k \ldots$ of depth-k bisimilarity by induction on $k < \omega$:

The relation $\leftrightarroweq^0 = M \times M'$ is the full relation. For every $k < \omega$, the relation $\leftrightarroweq^{k+1}$ links points that satisfy the base condition and the back and forth conditions with respect to $\leftrightarroweq^k$. That is, for every $\mathsf{m} \in M$ and $\mathsf{m}' \in M'$, $\mathsf{m} \leftrightarroweq^{k+1} \mathsf{m}'$ iff:

- (base) m and m' have the same color;
- (forth) for every $\mathsf{a} \in \mathsf{Act}$: whenever $\mathsf{m} \xrightarrow{\mathsf{a}}^{\mathcal{M}} \mathsf{n}$, there exists n' such that $\mathsf{m}' \xrightarrow{\mathsf{a}}^{\mathcal{M}'} \mathsf{n}'$ and $\mathsf{n} \leftrightarroweq^k \mathsf{n}'$;
- (back) for every $\mathsf{a} \in \mathsf{Act}$: whenever $\mathsf{m}' \xrightarrow{\mathsf{a}}^{\mathcal{M}'} \mathsf{n}'$, there exists n such that $\mathsf{m} \xrightarrow{\mathsf{a}}^{\mathcal{M}} \mathsf{n}$ and $\mathsf{n} \leftrightarroweq^k \mathsf{n}'$.

To characterize such $\leftrightarroweq^k$ in terms of a game, we enrich the game $\mathcal{G}_{\leftrightarroweq}(\mathcal{M}, \mathcal{M}')$ with a *counter* storing a natural number decremented along the play. The new game $\mathcal{G}_{\leftrightarroweq}^{<\omega}(\mathcal{M}, \mathcal{M}')$ is almost the same as $\mathcal{G}_{\leftrightarroweq}(\mathcal{M}, \mathcal{M}')$ except that an additional *countdown* step count, in which the counter will be decremented, is performed at the beginning of each round. In case the decrement is not possible because the counter has value 0, $\exists$ve wins.

That is, we extend the configurations of $\mathcal{G}_{\leftrightarroweq}(\mathcal{M}, \mathcal{M}')$ with counter values ranging over ω:

$$\begin{aligned}\mathsf{Conf}_{\exists} &= M \times M' \times \omega \times \{\mathsf{base}\} \ \sqcup \ M \times M' \times \omega \times \mathsf{Act} \times \{\mathsf{back}, \mathsf{forth}\} \\ \mathsf{Conf}_{\forall} &= M \times M' \times \omega \times \{\mathsf{b\&f}\} \ \sqcup \ M \times M' \times \omega \times \{\mathsf{count}\}\end{aligned}$$

and add a countdown step count when the counter is decremented:

1. In $(\mathsf{m}, \mathsf{m}', k, \mathsf{count})$, if $k = 0$ then there are no outgoing edges (meaning that $\forall$dam looses) and otherwise game moves deterministically to $(\mathsf{m}, \mathsf{m}', k-1, \mathsf{base})$
2. In $(\mathsf{m}, \mathsf{m}', k, \mathsf{base})$, if m and m' have different colors then there are no moves (so $\exists$ve looses immediately), otherwise the game moves deterministically to $(\mathsf{m}, \mathsf{m}', \mathsf{b\&f})$.
3. In $(\mathsf{m}, \mathsf{m}', \mathsf{b\&f})$, $\forall$dam chooses either:
   - $(\mathsf{n}, \mathsf{m}', k, \mathsf{a}, \mathsf{forth})$ such that $\mathsf{m} \xrightarrow{\mathsf{a}}^{\mathcal{M}} \mathsf{n}$ or
   - $(\mathsf{m}, \mathsf{n}', k, \mathsf{a}, \mathsf{back})$ such that $\mathsf{m}' \xrightarrow{\mathsf{a}}^{\mathcal{M}'} \mathsf{n}'$.
4. In $(\mathsf{n}, \mathsf{m}', k, \mathsf{a}, \mathsf{forth})$ or $(\mathsf{m}, \mathsf{n}', k, \mathsf{a}, \mathsf{back})$, $\exists$ve chooses $(\mathsf{n}, \mathsf{n}', k, \mathsf{count})$ such that $\mathsf{m}' \xrightarrow{\mathsf{a}}^{\mathcal{M}'} \mathsf{n}'$ or $\mathsf{m} \xrightarrow{\mathsf{a}}^{\mathcal{M}} \mathsf{n}$, respectively.

Since in each round the counter value decreases, there could be no infinite plays and so the above description is already complete. As promised, it captures depth-k bisimilarity:

$$\exists\text{ve wins } \mathcal{G}_{\leftrightarroweq}^{<\omega}(\mathcal{M}, \mathcal{M}'), (\mathsf{m}, \mathsf{m}', k, \mathsf{count}) \iff \mathcal{M}, \mathsf{m} \leftrightarroweq^k \mathcal{M}', \mathsf{m}'. \tag{2.3}$$

Proof. We proceed by induction on $k < \omega$. The case with $k = 0$ is immediate. Assuming (2.3) is true for some $k < \omega$, we will prove it for $k+1$. Observe that $\mathsf{m} \leftrightarroweq^{k+1} \mathsf{m}'$ iff both points satisfy (i) the base condition and (ii) the back and forth conditions with respect to $\leftrightarroweq^{k}$, meaning that for every $\xrightarrow{\mathsf{a}}$-successor of one of the points, there exists an $\xrightarrow{\mathsf{a}}$-successor of the other point such that the new points n and n' are k-step bisimilar. On the other hand, ∃ve wins $\mathcal{G}^{<\omega}_{\leftrightarroweq}(\mathcal{M}, \mathcal{M}'), (\mathsf{m}, \mathsf{m}', k+1, \mathsf{count})$ iff (i) m and m' have the same color and (ii) for every ∀dam's choice of an $\xrightarrow{\mathsf{a}}$-successor of one point, ∃ve can respond with an $\xrightarrow{\mathsf{a}}$-successor of the other point such that she wins from $(\mathsf{n}, \mathsf{n}', k, \mathsf{base})$ where n and n' are the new points. The first condition (i) is the same in both cases, whereas the second one (ii) is equivalent thanks to the induction hypothesis. □

Often, we will assume that configurations consist of a *position* v from some fixed set V called the *arena* and some extras such as information from some finite set, (natural- or ordinal-valued) counters, or a register storing a real value $a \in [0, 1]$. Technically, one could always define such arena V to be just the set of all configurations. However, the distinction between configurations and its underlying positions reflects different roles of the components of a configuration: it emphasizes similarities between various games that we are going to investigate and highlights finitary aspects of games with infinitely many configurations. We will call games where the set of configurations is equal to the arena *simple*.

For instance, the bisimulation game $\mathcal{G}_{\leftrightarroweq}(\mathcal{M}, \mathcal{M}')$ can be naturally seen as a simple game. On the contrary, we do *not* want to view the game $\mathcal{G}^{<\omega}_{\leftrightarroweq}(\mathcal{M}, \mathcal{M}')$ characterizing depth-k bisimilarity as a simple game. Instead, we think of configurations of $\mathcal{G}^{<\omega}_{\leftrightarroweq}(\mathcal{M}, \mathcal{M}')$ such as $(\mathsf{m}, \mathsf{m}', k, \mathsf{base})$ as consisting of a position $(\mathsf{m}, \mathsf{m}', \mathsf{base})$ and a counter value $k < \omega$. The two components are conceptually different and play different roles. This will be illustrated in Examples 2.2.10 and 4.1.3 where we derive the respective games $\mathcal{G}_{\leftrightarroweq}(\mathcal{M}, \mathcal{M}')$ and $\mathcal{G}^{<\omega}_{\leftrightarroweq}(\mathcal{M}, \mathcal{M}')$ from a more general theory.

2.2.1 Parity Winning Condition

Often, as with the bisimilarity game $\mathcal{G}_{\leftrightarroweq}(\mathcal{M}, \mathcal{M}')$ where ∃ve wins all the infinite plays, the winning condition of a game does not depend on finite prefixes of infinite plays but only on its limit properties. Given a winning condition $\mathsf{Win}_\exists \sqcup \mathsf{Win}_\forall$, we say that the condition is *prefix-independent* if for every infinite play $\pi\rho$ with a finite prefix π, $\pi\rho \in \mathsf{Win}_\exists$ iff $\rho \in \mathsf{Win}_\exists$.

A prefix-independent winning condition of great significance in logic and computer science is the *parity condition*. It is given by a *rank function* $\mathsf{rank} : \mathsf{Conf} \to \mathcal{R}$ that maps each configuration to an element of a fixed finite linear order $\mathcal{R} = \mathcal{R}_\exists \sqcup \mathcal{R}_\forall$ divided between players ∃ve and ∀dam. We say that player P *owns* rank $r \in \mathcal{R}$ if $r \in \mathcal{R}_P$. In case of an infinite play, the owner of the greatest rank r that appeared infinitely often *looses*.

The parity condition such that ∃ve looses all infinite plays is called a *reachability condition* and corresponds to $\mathcal{R} = \mathcal{R}_\exists$. Symmetrically, if ∃ve wins all infinite plays (which corresponds to $\mathcal{R} = \mathcal{R}_\forall$) we call it a *safety condition*. The general parity condition can be thought of as a result of nesting of these two simple ones. With a reachability condition ∃ve has to win in finitely many steps, so in this sense every position is bad for her. With a parity condition, a visit to a rank $r \in \mathcal{R}_\exists$ is bad for her *unless* the play visits a more important rank later.

We call a game prefix-independent, parity, safety, or reachability iff its winning condition is of such type. We will often make a convenient assumption that does not decrease generality: the lowest rank of a given game is never the most important one appearing infinitely often in a play. We will denote such lowest rank by 0 and since its ownership is irrelevant for determining who wins, we will leave it unspecified.

Let us recall a very useful fact: *configurational determinacy* (usually called *positional determinacy* in the literature [25]) of parity games. A strategy σ is called *configurational* if in order to determine the next move γ_{k+1} it only looks at the current configuration γ_k instead of the entire play $\gamma_1 ... \gamma_k$. That is, if $\sigma(\pi) = \gamma$ and π' ends with the same configuration as π then also $\sigma(\pi') = \gamma$. Parity games are configurationally determined [25]:

Proposition 2.2.1. *Assume a game with a parity condition. If player P wins the game then (s)he does so with a strategy that is configurational.*

From now on, unless stated otherwise, in parity games we only consider configurational strategies.

2.2.2 Partial Games and Gamemulations

In this subsection we introduce tools that will allow us to manipulate games in a precise manner. Although they are implicit in the usual approach to classical parity games and μ-ML (as found e.g. in [37]), a precise analysis now will pay off later when we apply it to more involved games.

Partial Games. Sometimes we will want to decompose games into smaller fragments. In what follows we will be mostly interested in analyzing ∃ve's strategies, so for the sake of simpler notation we focus on her, although analogous notions for ∀dam could be defined symmetrically.

Given a game $\mathcal{G}$ and a subset $\mathcal{S} \subseteq \mathsf{Conf}$ of its *stopping* configurations, the *partial game* $\mathcal{G}|\mathcal{S}$ is played the same as $\mathcal{G}$, except that upon moving to a configuration in $\mathcal{S}$ the play ends with a draw, meaning that no player looses or wins (note that *starting* in $\mathcal{S}$ does not count as a *move* to $\mathcal{S}$, so plays starting there are not necessarily stopped). For a non-loosing strategy σ for

$\exists$ve in $\mathcal{G}|\mathcal{S},\gamma$, the set of its *exit configurations*, denoted $\mathsf{exit}(\sigma) \subseteq \mathcal{S}$, consists of all the configurations $\delta \in \mathcal{S}$ such that some σ-play ends in δ.

Exit-better & Exit-equivalent. We think of a non-loosing $\exists$ve's strategy σ in $\mathcal{G}|\mathcal{S},\gamma$ as a candidate for a winning strategy in $\mathcal{G},\gamma$: $\exists$ve wins all σ-plays except for the ones that reach $\mathcal{S}$, in which case (s)he may need to continue with another strategy. In this light, σ' is better than σ if $\mathsf{exit}(\sigma') \subseteq \mathsf{exit}(\sigma)$, for there are fewer possible scenarios where $\exists$ve is not guaranteed to win. Since we want to compare different games, we will be interested in such comparisons mediated by a relation between configurations.

Assume another $\mathcal{G}'$ with configurations Conf', a subset $\mathcal{S}' \subseteq \mathsf{Conf}'$ and a relation $S \subseteq \mathcal{S} \times \mathcal{S}'$. To avoid overloaded notation, unless stated otherwise we assume that $\pi_1[S] = \mathcal{S}$ and $\pi_2[S] = \mathcal{S}'$ so that the sets $\mathcal{S}$ and $\mathcal{S}'$ are encoded by S. Given non-loosing strategies σ in $\mathcal{G}|\mathcal{S},\gamma$ and σ' in $\mathcal{G}'|\mathcal{S}',\gamma'$, we call σ' *exit-better with respect to* S than σ, and denote it $\sigma \lhd_S \sigma'$, iff:

$$\mathsf{exit}(\sigma') \subseteq S[\mathsf{exit}(\sigma)]$$

i.e. for every σ'-play stopped at some γ' there exists a σ-play stopping in γ with $\gamma S \gamma'$.

We also define the relation $\lhd_S$ between games: $\mathcal{G},\gamma \lhd_S \mathcal{G}',\gamma'$ means that for every non-loosing σ in $\mathcal{G}|\mathcal{S},\gamma$ there exists a non-loosing σ' in $\mathcal{G}'|\mathcal{S}',\gamma'$ such that $\sigma \lhd_S \sigma'$. The games $\mathcal{G},\gamma$ and $\mathcal{G}',\gamma'$ are *exit-equivalent* with respect to S, denoted $\mathcal{G},\gamma \bowtie_S \mathcal{G}',\gamma'$, if $\mathcal{G},\gamma \lhd_S \mathcal{G}',\gamma'$ and $\mathcal{G}',\gamma' \lhd_{S^{-1}} \mathcal{G},\gamma$. The relation of being exit-better composes, in the sense that $\mathcal{G},\gamma \lhd_S \mathcal{G}',\gamma'$ and $\mathcal{G}',\gamma' \lhd_{S'} \mathcal{G}'',\gamma''$ implies $\mathcal{G},\gamma \lhd_{S \circ S'} \mathcal{G}'',\gamma''$ (or, more succinctly, $\lhd_S \circ \lhd_{S'} \subseteq \lhd_{S \circ S'}$). The same is true for exit-equivalence.

In case $\mathcal{S} = \emptyset$, the game $\mathcal{G}|\emptyset$ is the same as $\mathcal{G}$, non-loosing strategies are actually winning and thus $\mathcal{G},\gamma \lhd_\emptyset \mathcal{G}',\gamma'$ means that if $\exists$ve wins $\mathcal{G},\gamma$ then she also wins $\mathcal{G}',\gamma'$. We hence denote $\lhd_\emptyset$ as $\lhd$ and $\bowtie_\emptyset$ as $\bowtie$.

The following proposition says that in the case of prefix-independent games exit-equivalence is compositional.

Proposition 2.2.2. *Assume two prefix-independent games $\mathcal{G},\gamma_I$ and $\mathcal{G}',\gamma'_I$ with pairs of disjoint subsets $\mathcal{S}_+ \sqcup \mathcal{S} \subseteq \mathsf{Conf}$ and $\mathcal{S}'_+ \sqcup \mathcal{S}' \subseteq \mathsf{Conf}'$ together with relations $S_+ \subseteq \mathcal{S}_+ \times \mathcal{S}'_+$ and $S \subseteq \mathcal{S} \times \mathcal{S}'$. If:*

$$\mathcal{G},\gamma_I \lhd_{S_+ \sqcup S} \mathcal{G}',\gamma'_I$$

and for all $(\gamma,\gamma') \in S_+$:

$$\mathcal{G},\gamma \lhd_S \mathcal{G}',\gamma',$$

then also:

$$\mathcal{G},\gamma_I \lhd_S \mathcal{G}',\gamma'_I.$$

Proof. We need to translate non-loosing $\exists$ve's strategies in $\mathcal{G}|\mathcal{S},\gamma_I$ to $\mathcal{G}'|\mathcal{S}',\gamma'_I$. Assume a strategy σ for $\mathcal{G}|\mathcal{S},\gamma_I$. We construct σ' for $\mathcal{G}'|\mathcal{S},\gamma'_I$ with $\mathsf{exit}(\sigma') \subseteq S[\mathsf{exit}(\sigma)]$. The key idea is to decompose σ into smaller parts, translate these parts to $\mathcal{G}'$ and obtain σ' by putting these translated pieces together. Both $\mathcal{G}|\mathcal{S},\gamma_I$ and $\mathcal{G}'|\mathcal{S}',\gamma'_I$ decompose into two phases: before and after the first move to $\mathcal{S}_+$ and $\mathcal{S}'_+$, respectively (if a play never enters configuration from $\mathcal{S}_+$ or $\mathcal{S}'_+$, the first phase may continue forever without moving to the second one). The new strategy σ' is as follows:

1. Since the first phase of $\mathcal{G}|\mathcal{S},\gamma_I$ ends with a move to $\mathcal{S}_+$, it can be identified with $\mathcal{G}|\mathcal{S}_+ \sqcup \mathcal{S},\gamma_I$. In this view, the strategy σ_0 being σ restricted to the first phase can be thought of as a non-loosing strategy for $\mathcal{G}|\mathcal{S}_+ \sqcup \mathcal{S},\gamma_I$. The assumption $\mathcal{G},\gamma_I \lhd_{S_+ \sqcup S} \mathcal{G},\gamma'_I$ implies existence of a non-loosing strategy σ'_0 for $\mathcal{G}'|\mathcal{S}'_+ \sqcup \mathcal{S}',\gamma'_I$ with:

   $$\mathsf{exit}(\sigma'_0) \subseteq (S_+ \sqcup S)[\mathsf{exit}(\sigma_0)].$$

   $\exists$ve plays according to such σ'_0 until a visit to $\mathcal{S}'_+ \sqcup \mathcal{S}'$. If this never happens she wins. Otherwise, disjointness of domains and codomains of S_+ and S imply:

   $$\mathsf{exit}(\sigma'_0) \cap \mathcal{S}'_+ \subseteq S_+[\mathsf{exit}(\sigma_0)] \qquad \text{and} \qquad \mathsf{exit}(\sigma'_0) \cap \mathcal{S}' \subseteq S[\mathsf{exit}(\sigma_0)].$$

   Thus, if the play reaches $\gamma' \in \mathcal{S}'$ then $\gamma S \gamma'$ for some $\gamma \in S[\mathsf{exit}(\sigma_0)] \subseteq S[\mathsf{exit}(\sigma)]$ and we are done.

2. The only remaining case is when after a play π' the game $\mathcal{G}'|\mathcal{S}'$ moves to $\gamma' \in \mathcal{S}'_+$. Prefix-independence of $\mathcal{G}'$ means that the history π' is irrelevant for winning the remaining game and it suffices if we provide $\exists$ve with a non-loosing strategy $\sigma_{\gamma'}$ in $\mathcal{G}'|\mathcal{S}',\gamma'$ with $\mathsf{exit}(\sigma_{\gamma'}) \subseteq S[\mathsf{exit}(\sigma)]$.

   Since $\gamma' \in \mathsf{exit}(\sigma'_0) \cap \mathcal{S}'_+$, we know that $\gamma' \in S_+[\mathsf{exit}(\sigma_0)]$. This means that $\gamma S_+ \gamma'$ for some $\gamma \in \mathcal{S}_+$ reachable from γ_I by a σ_0-play π (which, by definition of σ_0, is also a σ-play). Hence, the strategy σ_γ defined for every play ρ as:

   $$\sigma_\gamma(\rho) = \sigma(\pi\rho)$$

   is a non-loosing strategy for $\mathcal{G}|\mathcal{S},\gamma$ (which follows from prefix-independence of $\mathcal{G}$) with $\mathsf{exit}(\sigma_\gamma) \subseteq \mathsf{exit}(\sigma)$. Then, the assumption $\mathcal{G},\gamma \lhd_S \mathcal{G}',\gamma'$ allows to translate σ_γ to our desired $\sigma'_\gamma = \sigma_{\gamma'}$ in $\mathcal{G}'|\mathcal{S}',\gamma'$. Since the bigger X the bigger $S[X]$:

   $$\mathsf{exit}(\sigma'_\gamma) \subseteq S[\mathsf{exit}(\sigma_\gamma)] \subseteq S[\mathsf{exit}(\sigma)]$$

   which completes the proof.

$\square$

Gamemulations & Bisimulations. The connection between games and bisimilarity is twofold. We have already seen how bisimilarity can be captured in terms of a game, as expressed by (2.2). Let us now have a look at a complementary perspective: invariance of games under bisimilarity.

Any game can be seen as a monomodal model $\mathcal{M} = (M, \rightarrow^{\mathcal{M}}, \tau_{\exists}^{\mathcal{M}}, \tau_{\forall}^{\mathcal{M}})$ with the set of all configurations $\mathsf{Conf} = M$ as its points, the only accessibility relation interpreted as $\mathsf{Mov} = \rightarrow^{\mathcal{M}}$ and ownership of configurations encoded by two atomic propositions $\tau_{\exists}$ and $\tau_{\forall}$. We will show how a bisimulation between two games $\mathcal{G}, \gamma$ and $\mathcal{G}', \gamma'$ allows for a translation of strategies in one game to strategies in the other one. Under additional assumptions on the bisimulation, the translation of a winning strategy yields a strategy that is also winning.

For a more fine-grained analysis, we start with a notion weaker than bisimilarity that is not necessarily symmetric and only provides a one-way translation of strategies.

Definition 2.2.3. A a relation $G \subseteq \mathsf{Conf} \times \mathsf{Conf}'$ between sets of configurations of two (not necessarily distinct) games $\mathcal{G}$ and $\mathcal{G}'$ is called a *gamemulation* from $\mathcal{G}$ to $\mathcal{G}'$ if (i) it preserves the winning condition, meaning that for infinite plays π and π', $\pi G \pi'$ (pointwise) implies that π and π' have the same winner; (ii) for every $\gamma G \gamma'$, the following three conditions are satisfied:

- (base) γ and γ' have the same owner;
- (forth) whenever $\gamma \in \mathsf{Conf}_{\exists}$ and $\gamma \mathsf{Mov} \delta$, there exists $\delta' \in \mathsf{Conf}'$ such that $\gamma' \mathsf{Mov} \delta'$ and $\delta G \delta'$;
- (back) whenever $\gamma' \in \mathsf{Conf}'_{\forall}$ and $\gamma' \mathsf{Mov} \delta'$, there exists $\delta \in \mathsf{Conf}$ such that $\gamma \mathsf{Mov} \delta$ and $\delta G \delta'$.

R is a gamemulation from $\mathcal{G}, \gamma$ to $\mathcal{G}', \gamma'$, denoted $(\mathcal{G}, \gamma) G (\mathcal{G}', \gamma')$, if it is a gamemulation from $\mathcal{G}$ to $\mathcal{G}'$ such that $\gamma G \gamma'$. In case $\mathcal{G} = \mathcal{G}'$ and G is an order, we call it a *gamemulation order*.

The above definition is crafted specifically to enable translations of $\exists$ve's strategies.

Proposition 2.2.4. *Assume games $(\mathcal{G}, \gamma) G (\mathcal{G}', \gamma')$ linked by a gamemulation G. For every $\exists$ve's strategy σ in $\mathcal{G}, \gamma$ she has a σ' in $\mathcal{G}', \gamma'$ with the property that:*

- ▶ *for every σ'-play π' (finite or infinite),*
- ▶ *there exists a σ-play π with $\pi G \pi'$.*

In particular, for every $S \subseteq G$ we have $\sigma \lhd_S \sigma'$ and since σ is arbitrary it follows that $(\mathcal{G}, \gamma) \lhd_S (\mathcal{G}', \gamma')$.

Proof. The strategy σ' is constructed by induction on the length of plays, preserving the condition from the formulation of the proposition as an invariant. Assume that after a play π' the game has arrived at a configuration γ'. By the induction hypothesis, there exists π such that $\pi G \pi'$. In particular, π ends with γ that has the same owner as γ'. There are two cases:

- If γ and γ' belong to $\exists$ve, she looks at $\sigma(\pi) = \delta$. The forth condition implies existence of a legal move from γ' to δ' in $\mathcal{G}'$ and so we put $\sigma'(\pi') = \delta'$. Since $\delta G \delta'$, the invariant is preserved, for the σ-play $\pi\delta$ corresponds to the σ'-play $\pi'\delta'$.

- Symmetrically, if $\forall$dam owns both γ and γ' then the back condition implies that for every move to δ' in $\mathcal{G}'$ there is a legal (and hence extending π to a σ-play $\pi\delta$) move to δ with $\delta G \delta'$.

It follows from the construction of σ' that it has the desired property with respect to finite plays. The case with infinite ones follows from the observation that (i) an infinite play π' is consistent with σ' iff all its finite prefixes $\pi'_1, \pi'_2, ...$ are σ'-plays and (ii) given the corresponding σ-plays $\pi_1, \pi_2, ...$ from the invariant, π_i is a prefix of π_j whenever $i < j$. $\square$

Note that any relation $G \subseteq \mathsf{Conf} \times \mathsf{Conf}'$ between configurations of games $\mathcal{G}$ and $\mathcal{G}'$ that preserves the winning condition is a bisimulation iff both G and its inverse G^{-1} are gamemulations. As a corollary we get that games are invariant under bisimulations preserving the winning condition.

We call a subset $\mathcal{S} \subseteq \mathsf{Conf}$ of configurations of a (possibly partial) game $\mathcal{G}$ *victory-dominating* if the winner of any infinite play π visiting $\mathcal{S}$ infinitely often depends only on the subsequence $\pi \cap \mathcal{S}$ consisting of all π's elements from $\mathcal{S}$. Given another $\mathcal{G}'$ with $\mathcal{S}' \subseteq \mathsf{Conf}'$ we call a relation $S \subseteq \mathcal{S} \times \mathcal{S}'$ victory-dominating if $\mathcal{S}$ and $\mathcal{S}'$ are victory-dominating and whenever plays π and π' visit configurations from $\mathcal{S}$ and $\mathcal{S}'$ infinitely often: $(\pi \cap \mathcal{S}) S (\pi' \cap \mathcal{S}')$ (pointwise) implies that π and π' have the same winner. That is, if the i-th configuration from $\mathcal{S}$ in π is linked by S to the i-th configuration from $\mathcal{S}'$ in π', for all $i < \omega$, then π and π' have the same winner.

A typical example of this is with parity games. If r is the most important rank in $\mathcal{G}$ and all the elements of $\mathcal{S}$ have rank r then $\mathcal{S}$ is victory-dominating. Hence, given another $\mathcal{G}'$ with analogous r' and $\mathcal{S}'$, if r and r' have the same owner then every $S \subseteq \mathcal{S} \times \mathcal{S}'$ is victory-dominating. In case of partial games it suffices that such r and r' are the most important among configurations of $\mathcal{G}$ and $\mathcal{G}'$ *other than the stopping ones*, because no stopping configuration can appear in an infinite play. In the context of (partial) parity games we will often use this observation without explicitly mentioning it.

We now introduce an important lemma that allows to deduce equivalence of parity games from equivalence of its pieces.

Lemma 2.2.5 (Decomposition Lemma)**.** *Consider parity games $\mathcal{G}$ and $\mathcal{G}'$ with disjoint subsets $\mathcal{S}_+ \sqcup \mathcal{S} \subseteq \mathsf{Conf}$ and $\mathcal{S}'_+ \sqcup \mathcal{S}' \subseteq \mathsf{Conf}'$ and relations $S_+ \subseteq \mathcal{S}_+ \times \mathcal{S}'_+$ and $S \subseteq \mathcal{S} \times \mathcal{S}'$. Assume that S_+ is victory-dominating between $\mathcal{G}|\mathcal{S}$ and $\mathcal{G}'|\mathcal{S}'$. If:*

$$\mathcal{G}, \gamma \lhd_{S_+ \sqcup S} \mathcal{G}', \gamma'$$

for all $(\gamma, \gamma') \in S_+$ then also:

$$\mathcal{G}, \gamma \lhd_S \mathcal{G}', \gamma'.$$

for all $(\gamma, \gamma') \in S_+$.

Remark 2.2.6. We will usually us an immediate symmetric corollary of the above lemma where $\bowtie$ replaces $\lhd$ in both expressions and refer to it as *the* Decomposition Lemma.

Proof. As in the proof of Proposition 2.2.2, given a non-loosing σ in $\mathcal{G}|\mathcal{S}, \gamma$ we need to translate it to σ' in $\mathcal{G}'|\mathcal{S}', \gamma'$ so that $\mathsf{exit}(\sigma') \subseteq S[\mathsf{exit}(\sigma)]$. Again, we decompose the game $\mathcal{G}|\mathcal{S}$ into phases, look at the pieces of σ in these phases and translate each such piece to a corresponding phase of $\mathcal{G}'|\mathcal{S}'$. Then, the translated pieces of strategies compose into our desired σ' in $\mathcal{G}'|\mathcal{S}'$. The difference from the previous proof is that instead of splitting the game $\mathcal{G}|\mathcal{S}$ into two phases we split it into *infinitely* many ones, each consisting of plays between consecutive visits to $\mathcal{S}_+$; similarly with $\mathcal{G}'|\mathcal{S}'$ and $\mathcal{S}'_+$.

Since $\mathcal{G}$ is a parity game, we assume that σ is configurational. Thus, whenever the game arrives at some $\delta \in \mathcal{S}_+$ the behavior of σ depends only on δ and not the entire history of the play. It follows that σ can be presented as a set of strategies: for each $\delta \in \mathcal{S}_+$ a strategy σ_δ in $\mathcal{G}|\mathcal{S} \sqcup \mathcal{S}_+, \delta$ telling $\exists$ve how to play from δ until the next visit to $\mathcal{S}_+$. Then, the moves dictated by σ are precisely the ones given by σ_δ whenever δ is the last seen configuration from $\mathcal{S}_+$.

The assumption $\mathcal{G}, \delta \lhd_{S_+ \sqcup S} \mathcal{G}', \delta'$ enables a translation $\sigma_\delta \mapsto \sigma_{\delta'}$ for every $(\delta, \delta') \in S_+$ such that $\sigma_{\delta'}$ is non-loosing in $\mathcal{G}'|\mathcal{S}' \sqcup \mathcal{S}'_+, \delta'$ and $\mathsf{exit}(\sigma_{\delta'}) \subseteq (S \sqcup S_+)[\mathsf{exit}(\sigma_\delta)]$. Hence, $\exists$ve may play in $\mathcal{G}'|\mathcal{S}'$ by using $\sigma_{\delta'}$ whenever the last seen configuration from $\mathcal{S}'_+$ was δ'. Denote such strategy by σ'. It satisfies:

- ► For every σ'-play π' there exists a σ-play π with $(\pi \cap \mathcal{S}_+)\ S_+\ (\pi' \cap \mathcal{S}'_+)$.

The above can be checked by induction on the number of visits to $\mathcal{S}'_+$ and easily implies that σ' is indeed a legal, winning strategy in $\mathcal{G}'|\mathcal{S}'$ with all the required properties.

Instead of presenting the technical details of the above reasoning, which are in similar spirit to the proofs of Propositions 2.2.2 and 2.2.4, we will use the latter as a blackbox. Given the discussion above, we may view $\mathcal{G}|\mathcal{S}, \gamma$ as a sequence of alternating choices $\gamma_1\sigma_1\gamma_2\sigma_2...$: starting with $\gamma_1 = \gamma$, from γ_i:

1. $\exists$ve picks a positional σ_i in $\mathcal{G}|\mathcal{S}_+ \sqcup \mathcal{S}, \gamma_i$ and then

2. $\forall$dam chooses $\gamma_{i+1} \in \mathsf{exit}(\sigma_i)$; if $\gamma_{i+1} \in \mathcal{S}$ the game stops there, otherwise $\gamma_{i+1} \in \mathcal{S}_+$ and it continues from γ_{i+1}.

Denote such game by $\widehat{\mathcal{G}}|\mathcal{S}$. Formally, $\widehat{\mathcal{G}}$ has configurations:

$$\widehat{\mathsf{Conf}}_\exists = \mathcal{S}_+ \sqcup \mathcal{S} \qquad \text{and} \qquad \widehat{\mathsf{Conf}}_\forall = \bigcup_{\delta \in \mathcal{S}_+} \mathbb{S}_\delta$$

and moves:

$$\widehat{\mathsf{Mov}} = \{(\delta, \sigma) \mid \delta \in \mathcal{S}_+, \sigma \in \mathbb{S}_\delta\} \cup \{(\sigma, \delta) \mid \delta \in \mathsf{exit}(\sigma)\}$$

where:

$$\mathbb{S}_\delta = \{\sigma \mid \sigma \text{ is a non-loosing strategy in } \mathcal{G}|\mathcal{S}_+ \sqcup \mathcal{S}, \delta\}.$$

Towards the winning condition note that every infinite (and hence not containing elements of $\mathcal{S}$) play $\widehat{\pi}$ in $\widehat{\mathcal{G}}$ can be extended to (at least one) π in $\mathcal{G}$ satisfying $\pi \cap \mathcal{S}_+ = \widehat{\pi}$. Moreover, the winner of π depends only on $\widehat{\pi}$. We therefore define:

$$\widehat{\mathsf{Win}}_\exists = \{\widehat{\pi} \mid \exists\text{ve wins some (equivalently: all) } \pi \text{ in } \mathcal{G} \text{ with } \pi \cap \mathcal{S}_+ = \widehat{\pi}\}$$

and $\widehat{\mathsf{Win}}_\forall$ as its complement. The definition of $\widehat{\mathcal{G}}$ is tailored so that the games $\mathcal{G}|\mathcal{S}$ and $\widehat{\mathcal{G}}|\mathcal{S}$ are equivalent:

$$\mathcal{G}, \gamma \bowtie_{\mathsf{Id}(\mathcal{S})} \widehat{\mathcal{G}}, \gamma.$$

The game $\widehat{\mathcal{G}'}$ is defined analogously. The assumptions of the lemma imply that the relation:

$$S \sqcup \{(\sigma, \sigma') \mid \mathsf{exit}(\sigma') \subseteq (S_+ \sqcup S)[\mathsf{exit}(\sigma)]\}$$

is a gamemulation from $\widehat{\mathcal{G}}$ to $\widehat{\mathcal{G}'}$. Since it includes S, Proposition 2.2.4 entails the middle relation in:

$$\mathcal{G}, \gamma \lhd_{\mathsf{Id}(\mathcal{S})} \widehat{\mathcal{G}}, \gamma \lhd_S \widehat{\mathcal{G}'}, \gamma' \lhd_{\mathsf{Id}(\mathcal{S}')} \mathcal{G}', \gamma'$$

and since $\mathsf{Id}(\mathcal{S}) \circ S \circ \mathsf{Id}(\mathcal{S}') = S$ we get:

$$\mathcal{G}, \gamma \lhd_S \mathcal{G}', \gamma'$$

which proves the lemma. $\square$

2.2.3 Games for Fixpoints

Let us recall the game-theoretic characterization of fixpoints. Assume a set X and a monotone operation $f : \mathcal{P}(X) \to \mathcal{P}(X)$ on its powerset $\mathcal{P}(X)$ seen as a complete lattice ordered by set inclusion.

Definition 2.2.7. We define the *least fixpoint game* $\mathcal{G}_\mu(f)$ and the *greatest fixpoint game* $\mathcal{G}_\nu(f)$. Both games are played in two-step rounds as follows:

1. from a position $x \in X$, $\exists$ve chooses $Y \in \mathcal{P}(X)$ with $x \in f(Y)$,

2. $\forall$dam chooses $y \in Y$ and the next round starts from y.

Formally, both games have configurations:

$$\mathsf{Conf}_\exists = X \quad \text{and} \quad \mathsf{Conf}_\forall = \mathcal{P}(X)$$

and moves:

$$\mathsf{Mov} = \{(x, Y) \mid x \in f(Y)\} \cup \{(Y, x) \mid x \in Y\}.$$

The only difference between μ and ν is in the winning condition: in case of an infinite play $\exists$ve looses in $\mathcal{G}_\mu(f)$ and wins in $\mathcal{G}_\nu(f)$.

The games characterize fixpoints of f as follows.

Theorem 2.2.8. *For every $x \in X$:*

1. *$\exists$ve wins $\mathcal{G}_\mu(f)$ from x iff $x \in \mathsf{LFP}.f$.*

2. *$\exists$ve wins $\mathcal{G}_\nu(f)$ from x iff $x \in \mathsf{GFP}.f$.*

Proof. We start with the following proposition:

Proposition 2.2.9. *For every ordinal α and element $x \in X$:*

$$x \in f_\mu^\alpha \quad \iff \quad \exists_{\beta<\alpha}.\exists_Y.\ x \in f(Y) \quad \textit{and} \quad Y \subseteq f_\mu^\beta \tag{2.4}$$
$$x \in f_\nu^\alpha \quad \iff \quad \forall_{\beta<\alpha}.\exists_Y.\ x \in f(Y) \quad \textit{and} \quad Y \subseteq f_\nu^\beta \tag{2.5}$$

Proof. For the first equivalence (2.4), recall the definition:

$$f_\mu^\alpha = \bigcup_{\beta<\alpha} f(f_\mu^\beta).$$

If $x \in f_\mu^\alpha$, then there exists $\beta < \alpha$ such that $x \in f(f_\mu^\beta)$ and hence $Y = f_\mu^\beta$ witnesses that the claim is true. Conversely, if $x \notin f_\mu^\alpha$ then for every $\beta < \alpha$ and Y such that $x \in f(Y)$ we have $Y \not\subseteq f_\mu^\beta$. This is because otherwise monotonicity would imply $x \in f(Y) \subseteq f(f_\mu^\beta) \subseteq f_\mu^\alpha$ which is impossible.

Towards the second equivalence (2.5), recall that:

$$f_\nu^\alpha = \bigcap_{\beta<\alpha} f(f_\nu^\beta).$$

When $x \in f_\nu^\alpha$ then for every $\beta < \alpha$ we have $x \in f(f_\nu^\beta)$ and hence $Y = f_\nu^\beta$ witnesses the claim for β. If $x \notin f_\nu^\alpha$, then there exists $\beta < \alpha$ with $x \notin f(f_\nu^\beta)$ and hence by monotonicity $x \in f(Y)$ implies $Y \not\subseteq f_\nu^\beta$. □

Let us prove Theorem 2.2.8. For the first item, take α big enough so that $f_\mu^\alpha = f_\mu^{\alpha+1} = \mathsf{LFP}.f$. If $x \notin \mathsf{LFP}.f = f_\mu^{\alpha+1}$ then by (2.4) for every $\beta < \alpha + 1$ and a legal move Y for ∃ve we have $Y \not\subseteq f_\mu^\beta$. In particular, taking $\beta = \alpha$ we get that every legal Y contains an element $y \notin f_\mu^\alpha = \mathsf{LFP}.f$. Thus, ∀dam can keep picking elements from outside of $\mathsf{LFP}.f$ and never get stuck. Since $\mathcal{G}_\mu(f)$ is a safety game such strategy guarantees him victory.

Conversely, whenever $y \in f_\mu^\beta$ then (2.4) implies existence of a set Y and an ordinal $\beta' < \beta$ such that Y is a legal move from y and $Y \subseteq f_\mu^{\beta'}$. If ∃ve chooses such Y, then every response $z \in Y$ chosen by ∀dam will belong to $f_\mu^{\beta'}$. Hence, starting from x ∃ve can keep picking sets $Y_0 \supseteq Y_1 \supseteq ...$ with corresponding ordinals $\alpha > \beta_0 > \beta_1 > ...$ such that $Y_i \subseteq f_\mu^{\beta_i}$. By well-foundedness of Ord, this must end in finitely many steps and so ∃ve wins.

For the second item of the theorem, take α for which $f_\nu^\alpha = f_\nu^{\alpha+1} = \mathsf{GFP}.f$. ∃ve can win from $x \in \mathsf{GFP}.f$ if she plays maintaining as an invariant that all her configurations that are visited belong to $\mathsf{GFP}.f$. This is possible because whenever $y \in \mathsf{GFP}.f = f_\nu^{\alpha+1}$ then (2.5) implies existence of $Y \subseteq f_\nu^\alpha$ with $y \in f(Y)$. Such Y is a legal move for ∃ve from y and all possible ∀dam's responses belong to $f_\nu^\alpha = \mathsf{GFP}.f$.

For the last implication, assume that $x \notin f_\nu^\alpha$. Whenever $y \notin f_\nu^\beta$, (2.5) implies existence of $\beta' < \beta$ such that for every Y that ∃ve may legally choose in y, this Y must contain an element $z \notin f_\nu^{\beta'}$. Thus, starting from x ∀dam can pick elements $y_0, y_1 ...$ with corresponding ordinals $\alpha > \beta_0 > \beta_1 > ...$ such that $y_i \notin f_\nu^{\beta_i}$. Again, well-foundedness of Ord implies that such process must terminate after finitely many rounds and so ∀dam will win. □

Example 2.2.10. Let us demonstrate the somewhat abstract Theorem 2.2.8 in action by showing how the characterization (2.2) of bisimilarity $\leftrightarroweq$ in terms of the bisimilarity game $\mathcal{G}_{\leftrightarroweq}(\mathcal{M}, \mathcal{M}')$ can be seen as its instance.

In order to apply the characterization, we need to view the notion of bisimilarity between $\mathcal{M}$ and $\mathcal{M}'$ as a fixpoint of some function. Consider the following operation $\mathsf{BIS} : \mathcal{P}(M \times M') \to \mathcal{P}(M \times M')$ on binary relations between the models:

$$\begin{aligned}\mathsf{BIS}(Z) = \{(\mathsf{m},\mathsf{m}') \mid\ & \forall_{\tau\in\mathsf{Prop}}.\ \mathsf{m}\in\tau^{\mathcal{M}} \iff \mathsf{m}'\in\tau^{\mathcal{M}'} \ \wedge \\ & \forall_{\mathsf{a}\in\mathsf{Act}}.\forall_{\mathsf{m}\xrightarrow{\mathsf{a}}^{\mathcal{M}}\mathsf{n}}.\exists_{\mathsf{m}'\xrightarrow{\mathsf{a}}^{\mathcal{M}'}\mathsf{n}'}.\ \mathsf{n}Z\mathsf{n}' \ \wedge \\ & \forall_{\mathsf{a}\in\mathsf{Act}}.\forall_{\mathsf{m}'\xrightarrow{\mathsf{a}}^{\mathcal{M}}\mathsf{n}'}.\exists_{\mathsf{m}\xrightarrow{\mathsf{a}}^{\mathcal{M}'}\mathsf{n}}.\ \mathsf{n}Z\mathsf{n}'\}.\end{aligned}$$

The three clauses in the above definition correspond to the base, forth and back conditions (with respect to Z) of a bisimulation relation, respectively. It follows that every relation Z between M and M' is a bisimulation iff $Z \subseteq \mathsf{BIS}(Z)$. Since the relation $\leftrightarroweq$ of bisimilarity is the greatest bisimulation, it is the greatest such relation and thus it equals to the greatest fixpoint of BIS:

$$\leftrightarroweq \ = \ \mathsf{GFP.BIS}.$$

According to Theorem 2.2.8, this means that $\exists$ve wins the game $\mathcal{G}_\nu(\mathsf{BIS})$ from (m,m') iff $\mathsf{m} \leftrightarroweq \mathsf{m}'$. Hence, to obtain (2.2) it suffices to show that the games $\mathcal{G}_\nu(\mathsf{BIS})$ and $\mathcal{G}_{\leftrightarroweq}(\mathcal{M},\mathcal{M}')$ have the same winner:

$$\mathcal{G}_\nu(\mathsf{BIS}),(\mathsf{m},\mathsf{m}') \bowtie \mathcal{G}_{\leftrightarroweq}(\mathcal{M},\mathcal{M}'),(\mathsf{m},\mathsf{m}',\mathsf{base}) \tag{2.6}$$

for all $\mathsf{m}\in M$ and $\mathsf{m}'\in M'$. Let us prove this equivalence.

Proof. Both games $\mathcal{G}_\nu(\mathsf{BIS})$ and $\mathcal{G}_{\leftrightarroweq}(\mathcal{M},\mathcal{M}')$ can be thought of as moving from one pair of points to another in rounds. Unravelling the definition of $\mathcal{G}_\nu(\mathsf{BIS})$ we get that from every (m,m'):

1. $\exists$ve picks a relation $Z \subseteq M\times M'$ such that $(\mathsf{m},\mathsf{m}')\in\mathsf{BIS}(Z)$, and then
2. $\forall$dam chooses $(\mathsf{n},\mathsf{n}')\in Z$ from which the next round starts.

On the other hand, in $\mathcal{G}_{\leftrightarroweq}(\mathcal{M},\mathcal{M}')$ from $(\mathsf{m},\mathsf{m}',\mathsf{base})$:

1. equivalence of atomic propositions is checked, then
2. $\forall$dam chooses $\mathsf{m}\xrightarrow{\mathsf{a}}\mathsf{n}$ or $\mathsf{m}'\xrightarrow{\mathsf{a}}\mathsf{n}'$ and
3. $\exists$ve responds with $\mathsf{m}'\xrightarrow{\mathsf{a}}\mathsf{n}'$ or $\mathsf{m}\xrightarrow{\mathsf{a}}\mathsf{n}$, respectively, so that the next position is $(\mathsf{n},\mathsf{n}',\mathsf{base})$.

Denote:

$$\mathcal{G} = \mathcal{G}_\nu(\mathsf{BIS}) \qquad \text{and} \qquad \mathcal{G}' = \mathcal{G}_{\leftrightarroweq}(\mathcal{M},\mathcal{M}').$$

We will translate between $\exists$ve's strategies in a round of either of the two games. Roughly, the translated strategy will allow to end the round in the same configurations as the original strategy, up to identifying every (n,n') in $\mathcal{G}$ with $(\mathsf{n},\mathsf{n}',\mathsf{base})$ in $\mathcal{G}'$.

Formally, we apply the Decomposition Lemma 2.2.5. To that end, we view both $\mathcal{G}$ and $\mathcal{G}'$ as parity games: we assign ranks r and r', both belonging to ∀dam, to all configurations. This way, ∀dam looses all infinite plays and hence $\mathcal{G}$ and $\mathcal{G}'$ are safety games, as desired. Denote sets of configurations:

$$\mathcal{S}_+ = M \times M' \qquad \text{and} \qquad \mathcal{S}'_+ = M \times M' \times \{\mathsf{base}\}$$

and define the relation:

$$S_+ = \{((\mathsf{m}, \mathsf{m}'), (\mathsf{m}, \mathsf{m}', \mathsf{base}))|\ \mathsf{m} \in M, \mathsf{m}' \in M'\}.$$

Substituting $S = \emptyset$ we conclude from the lemma that to prove (2.6) it suffices to show:

$$\mathcal{G}, (\mathsf{m}, \mathsf{m}') \bowtie_{S_+} \mathcal{G}', (\mathsf{m}, \mathsf{m}', \mathsf{base}) \tag{2.7}$$

for all $\mathsf{m} \in M$ and $\mathsf{m}' \in M'$. That is, we need to translate strategies between partial games $\mathcal{G}|\mathcal{S}_+, (\mathsf{m}, \mathsf{m}')$ and $\mathcal{G}'|\mathcal{S}'_+, (\mathsf{m}, \mathsf{m}', \mathsf{base})$ for every $\mathsf{m} \in M$ and $\mathsf{m}' \in M'$. We assume that m and m' satisfy the same atomic propositions, as otherwise ∃ve looses both games immediately and the equivalence becomes trivial.

Assume an ∃ve's strategy σ' in $\mathcal{G}'|\mathcal{S}'_+, (\mathsf{m}, \mathsf{m}', \mathsf{base})$. Such σ' gives her a reply $\mathsf{m}' \xrightarrow{\mathsf{a}} \mathsf{n}'$ to every $\mathsf{m} \xrightarrow{\mathsf{a}} \mathsf{n}$ (and symmetrically $\mathsf{m} \xrightarrow{\mathsf{a}} \mathsf{n}$ for every $\mathsf{m}' \xrightarrow{\mathsf{a}} \mathsf{n}'$). This can be summarized as a pair of functions $f_\mathsf{a} : Y_\mathsf{a} \to Y'_\mathsf{a}$ and $f'_\mathsf{a} : Y'_\mathsf{a} \to Y_\mathsf{a}$ for each a, where Y_a and Y'_a denote the $\xrightarrow{\mathsf{a}}$-successors of m and m', respectively. Consider the union:

$$Z = \bigcup_{\mathsf{a} \in \mathsf{Act}} \{(\mathsf{n}, f_\mathsf{a}(\mathsf{n}))\ |\ \mathsf{n} \in Y_\mathsf{a}\} \cup \{(f'_\mathsf{a}(\mathsf{n}), \mathsf{n})\ |\ \mathsf{n} \in Y'_\mathsf{a}\}$$

of the graphs of all f_a's and f'_a's. It follows that $(\mathsf{m}, \mathsf{m}') \in \mathsf{BIS}(Z)$, meaning that Z is a legal move in $\mathcal{G}_\nu(\mathsf{BIS})$. Let σ be the strategy that chooses this Z. Since $(\mathsf{n}, \mathsf{n}') \in Z$ iff some σ'-play leads to $(\mathsf{n}, \mathsf{n}', \mathsf{base})$ and $\mathsf{exit}(\sigma) = Z$:

$$\mathsf{exit}(\sigma) = S_+^{-1}[\mathsf{exit}(\sigma')]$$

as desired.

For the opposite translation, assume that ∃ve has a strategy σ in $\mathcal{G}|\mathcal{S}_+$ that legally picks a relation Z. Observe that whether $(\mathsf{m}, \mathsf{m}')$ belongs to $\mathsf{BIS}(Z)$ only depends on the part $Z \cap (Y \times Y')$ of Z between children Y and Y' of m and m', respectively. Hence:

$$(\mathsf{m}, \mathsf{m}') \in \mathsf{BIS}(Z \cap (Y \times Y')).$$

By definition of BIS, this means that for every $\xrightarrow{\mathsf{a}}$-successor n of m or n' of m' there is at least one $\xrightarrow{\mathsf{a}}$-successor n' or n of the other point, respectively, with $\mathsf{n}Z\mathsf{n}'$. This gives ∃ve a strategy σ', as she may respond with any such n' or n to every n or n', respectively. Similarly to the previous case $\mathsf{exit}(\sigma) = Z$ and so:

$$\mathsf{exit}(\sigma') \subseteq S_+[\mathsf{exit}(\sigma)]$$

which completes the proof of (2.7) and therefore also (2.6). □

2.3 Modal Logic

A logic that is of particular interest in this thesis is the *modal logic* ML. Fix sets of atomic propositions Prop and actions Act. Sentences of modal logic are given by the grammar:

$$\varphi ::= \top \mid \bot \mid \tau \mid \neg\tau \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \langle \mathsf{a} \rangle \varphi \mid [\mathsf{a}]\varphi$$

where $\tau \in \mathsf{Prop}$ and $\mathsf{a} \in \mathsf{Act}$. The two constructors $\langle \mathsf{a} \rangle$ and $[\mathsf{a}]$ are called *modal operators*. We call the set:

$$\mathsf{Lit} = \{\tau, \neg\tau \mid \tau \in \mathsf{Prop}\}$$

the *literals* over Prop. In the *monomodal* case, i.e. with $|\mathsf{Act}| = 1$, we skip the labels and write $\Diamond$ and $\Box$ instead of $\langle \mathsf{a} \rangle$ and $[\mathsf{a}]$.

We denote the set of all subformulae of a formula φ by $\mathsf{SubFor}(\varphi)$, where isomorphic subformulae are *not* identified. Hence, formally a subformula of φ is a node in its syntactic tree. We use this convention consistently for ML as well as for all its extensions considered in the thesis. Although irrelevant now, it will gain importance later and allow us to avoid unnecessarily convoluted definitions. Nonetheless, we will often abuse terminology and use the term subformula in the more usual sense: a *formula* whose syntactic tree is a subtree of the syntactic tree of φ.

Modal logic is interpreted in points of Kripke models. We always assume that the modal signature matches with the logic: it consists of unary and binary relational symbols Prop and $\{\xrightarrow{\mathsf{a}} \mid \mathsf{a} \in \mathsf{Act}\}$. Given such a model $\mathcal{M}$, the *semantics* $[\![\varphi]\!] \subseteq M$ is defined inductively as follows:

$$\begin{aligned}
[\![\top]\!] &= M \quad \text{and} \quad [\![\bot]\!] = \emptyset \\
[\![\tau]\!] &= \tau^{\mathcal{M}} \quad \text{and} \quad [\![\neg\tau]\!] = M - \tau^{\mathcal{M}} \\
[\![\varphi_1 \vee \varphi_2]\!] &= [\![\varphi_1]\!] \cup [\![\varphi_2]\!] \quad \text{and} \quad [\![\varphi_1 \wedge \varphi_2]\!] = [\![\varphi_1]\!] \cap [\![\varphi_2]\!] \\
[\![\langle \mathsf{a} \rangle\varphi]\!] &= \{\mathsf{m} \in M \mid \exists_{\mathsf{m} \xrightarrow{\mathsf{a}} \mathsf{n}}\ \mathsf{n} \in [\![\varphi]\!]\} \quad \text{and} \quad [\![[\mathsf{a}]\varphi]\!] = \{\mathsf{m} \in M \mid \forall_{\mathsf{m} \xrightarrow{\mathsf{a}} \mathsf{n}}\ \mathsf{n} \in [\![\varphi]\!]\}.
\end{aligned}$$

We say that $\mathcal{M}, \mathsf{m}$ (or just m if $\mathcal{M}$ is clear from the context) *satisfies* φ if $\mathsf{m} \in [\![\varphi]\!]$ and denote it by $\mathcal{M}, \mathsf{m} \models \varphi$. The *language* of a sentence φ is the class of all pointed models that satisfy φ. Given another $\mathcal{M}', \mathsf{m}'$ we call the models *modally equivalent* and denote it $\mathcal{M}, \mathsf{m} \equiv_{\mathsf{ML}} \mathcal{M}', \mathsf{m}'$ if they satisfy the same modal formulae.

Negation. Note that the negation $\neg$ is only allowed on the atomic propositions τ and not on arbitrary formulae. However, using de Morgan laws and the dualities between modal operators:

$$\neg\langle \mathsf{a} \rangle\neg\varphi \iff [\mathsf{a}]\varphi \quad \text{and} \quad \neg[\mathsf{a}]\neg\varphi \iff \langle \mathsf{a} \rangle\varphi,$$

every such negation could be pushed to the atomic propositions. The assumption that all formulae are already in such *negation normal form* will be convenient. Nevertheless, we will use connectives $\neg$, $\implies$ and $\iff$ as shorthands for their semantic equivalents.

Atomic Propositions vs. Colors. The set Prop of atomic propositions can be infinite. In that case it is not possible to describe the color of a point with a single formula. However, if the set Prop is finite then specifications in terms of colors and atomic propositions are equivalent. A color $c \in \mathcal{P}(\mathsf{Prop})$ can be described by a conjunction of all the literals consistent with it. Dually, a literal τ or $\neg\tau$ is equivalent to the disjunction of all colors consistent with that literal. Therefore, in such finitary case we will sometimes consider an equivalent syntax of ML with colors in place of literals.

Invariance under Bisimulation. A basic fact about modal logic is that it is *invariant under bisimulation*, meaning that it cannot distinguish two bisimilar points:

$$\mathcal{M}, \mathsf{m} \leftrightarroweq \mathcal{M}', \mathsf{m}' \qquad \implies \qquad \mathcal{M}, \mathsf{m} \equiv_{\mathsf{ML}} \mathcal{M}', \mathsf{m}'$$

for every $\mathcal{M}, \mathsf{m}$ and $\mathcal{M}', \mathsf{m}'$. The above fact follows from its refined version. Define the *modal-propositional depth* of a formula φ to be the maximal nesting of modal operators and atomic propositions in φ. For $P \subseteq \mathsf{Prop}$ and $k < \omega$ denote $\mathcal{M}, \mathsf{m} \equiv_{\mathsf{ML}}^{k,P} \mathcal{M}', \mathsf{m}'$ if $\mathcal{M}, \mathsf{m}$ and $\mathcal{M}', \mathsf{m}'$ satisfy the same formulae of modal-propositional depth at most k and only using atomic propositions from P (we skip the index P in the case $P = \mathsf{Prop}$). Since every formula uses a finite number of atomic propositions, checking modal equivalence boils down to checking it with respect to every finite subset of Prop:

$$\begin{gathered} \mathcal{M}, \mathsf{m} \equiv_{\mathsf{ML}}^{k} \mathcal{M}', \mathsf{m}' \\ \iff \\ \mathcal{M}, \mathsf{m} \equiv_{\mathsf{ML}}^{k,P} \mathcal{M}', \mathsf{m}' \quad \text{for every finite } P \subseteq \mathsf{Prop} \end{gathered}$$

for every $k < \omega$. Analogously, the relation $\leftrightarroweq^{k,P}$ is defined the same as $\leftrightarroweq^{k}$ except that the base condition is only checked against propositions in P. An easy induction on k shows that:

$$\mathcal{M}, \mathsf{m} \leftrightarroweq^{k,P} \mathcal{M}', \mathsf{m}' \qquad \iff \qquad \mathcal{M}, \mathsf{m} \equiv_{\mathsf{ML}}^{k,P} \mathcal{M}', \mathsf{m}'$$

for every *finite* $P \subseteq \mathsf{Prop}$ and $k < \omega$. This leads to a common proof technique: to show $\mathcal{M}, \mathsf{m} \equiv_{\mathsf{ML}} \mathcal{M}', \mathsf{m}'$ it suffices to prove $\mathcal{M}, \mathsf{m} \equiv_{\mathsf{ML}}^{k,P} \mathcal{M}', \mathsf{m}'$ for every $k < \omega$ and finite $P \subseteq \mathsf{Prop}$, and the latter one is equivalent to $\mathcal{M}, \mathsf{m} \leftrightarroweq^{k,P} \mathcal{M}', \mathsf{m}'$.

In particular, since $\leftrightarroweq^{k}$ implies $\leftrightarroweq^{k,P}$ for all $P \subseteq \mathsf{Prop}$ we get that $\leftrightarroweq^{k}$ implies $\equiv_{\mathsf{ML}}^{k}$ and consequently $\leftrightarroweq$ implies $\equiv_{\mathsf{ML}}$. In the special case when the set Prop is finite, the relations $\leftrightarroweq^{k}$ and $\equiv_{\mathsf{ML}}^{k}$ coincide.

In general, the relation $\mathcal{M}, \mathsf{m} \leftrightarroweq^{k,P} \mathcal{M}', \mathsf{m}'$ can be equivalently defined as $\mathcal{M}_0, \mathsf{m} \leftrightarroweq^{k} \mathcal{M}_0', \mathsf{m}'$ where $\mathcal{M}_0$ and $\mathcal{N}_0'$ are the respective reducts of $\mathcal{M}$ and $\mathcal{M}'$ (i.e. models for smaller vocabulary obtained by forgetting all the interpretations of atomic propositions outside of P). Thus, it follows from the game characterization (2.3) of $\leftrightarroweq^k$ that:

$$\exists\text{ve wins } \mathcal{G}_{\leftrightarroweq}^{<\omega,P}(\mathcal{M}, \mathcal{M}'), (\mathsf{m}, \mathsf{m}', k, \mathsf{count}) \iff \mathcal{M}, \mathsf{m} \leftrightarroweq^{k,P} \mathcal{M}', \mathsf{m}'. \quad (2.8)$$

where $\mathcal{G}_{\leftrightarroweq}^{<\omega,P}(\mathcal{M}, \mathcal{M}')$ is a modified version of the game $\mathcal{G}_{\leftrightarroweq}^{<\omega}(\mathcal{M}, \mathcal{M}')$ in which only the atomic propositions from P are checked in the base step.

Game Semantics for ML. Similarly to the notion of bisimilarity, the semantics of modal formulae can be naturally captured in terms of a game. Fix a model $\mathcal{M}$. Given a modal formula φ, the *semantic game* $\mathcal{G}(\varphi)$ is defined as follows. The set $\mathsf{Conf} = M \times \mathsf{SubFor}(\varphi)$ of configurations consists of points of $\mathcal{M}$ and subformulae of φ. The possible moves depend on the topmost connective and reflect the inductive definition of the semantics. In configurations:

- $(\mathsf{m}, \top)$ or $(\mathsf{m}, \bot)$, $\exists$ve immediately wins or looses, respectively;
- (m, τ) or $(\mathsf{m}, \neg\tau)$ $\exists$ve wins iff $\mathsf{m} \in \tau^{\mathcal{M}}$ or $\mathsf{m} \notin \tau^{\mathcal{M}}$, respectively;
- $(\mathsf{m}, \psi_1 \vee \psi_2)$ or $(\mathsf{m}, \psi_1 \wedge \psi_2)$, $\exists$ve or $\forall$dam, respectively, chooses $i \in \{1, 2\}$ and the game continues from (m, ψ_i);
- $(\mathsf{m}, \langle \mathsf{a} \rangle \psi)$ or $(\mathsf{m}, [\mathsf{a}]\psi)$, $\exists$ve or $\forall$dam, respectively, chooses an $\xrightarrow{\mathsf{a}}$-successor n of m and the game continues from (n, ψ).

Since at each step the formula component of a configurations is a strict subformula of the previous one, the game must end in at most $|\mathsf{SubFor}(\varphi)|$ many such steps. In particular, there are no infinite plays, and so the description above is already complete. As mentioned, the game characterizes the meaning of formulae.

Theorem 2.3.1 (ML Adequacy). *For all* $\mathsf{m} \in M$ *and* $\varphi \in \mathsf{ML}$*:*

$$\mathsf{m} \in [\![\varphi]\!] \iff \exists\textit{ve wins } \mathcal{G}(\varphi) \textit{ from } (\mathsf{m}, \varphi).$$

Proof. The above equivalence can be verified by a straightforward induction on the formula, so we only consider the case of a formula of shape $\langle \mathsf{a} \rangle \varphi$ for a demonstration. By definition, $\mathsf{m} \in [\![\langle \mathsf{a} \rangle \varphi]\!]$ iff there exists an n such that $\mathsf{m} \xrightarrow{\mathsf{a}} \mathsf{n}$ and $\mathsf{n} \in [\![\varphi]\!]$. On the other hand, $\exists$ve wins $\mathcal{G}(\langle \mathsf{a} \rangle \varphi)$ from $(\mathsf{m}, \langle \mathsf{a} \rangle \varphi)$ iff she can choose an $\xrightarrow{\mathsf{a}}$-successor n of m such that she wins from (n, φ) in $\mathcal{G}(\langle \mathsf{a} \rangle \varphi)$. By induction hypothesis, we know that for every $\mathsf{n} \in \mathcal{M}$, $\exists$ve wins from (n, φ) in $\mathcal{G}(\varphi)$ iff $\mathsf{n} \in [\![\varphi]\!]$. But the games $\mathcal{G}(\langle \mathsf{a} \rangle \varphi), (\mathsf{n}, \varphi)$ and $\mathcal{G}(\varphi), (\mathsf{n}, \varphi)$ are equivalent (since the only difference between them is the presence of configurations of shape $(\mathsf{m}, \langle \mathsf{a} \rangle \varphi)$ in the first game, but such configurations are not reachable by any play). Thus, it follows that ∃ve wins $\mathcal{G}(\langle \mathsf{a} \rangle \varphi)$ from $(\mathsf{m}, \langle \mathsf{a} \rangle \varphi)$ iff $\mathsf{m} \in [\![\langle \mathsf{a} \rangle \varphi]\!]$, as desired. $\square$

Apart from modal logic, we will refer to first-order logic FO and assume the reader to be familiar with its basics ([24] is a good reference). Occasionally, we will also mention monadic second-order logic MSO, an extension of FO with quantification over sets of elements (i.e. unary relations). Although MSO has proven important in computer science, knowing it is not a prerequisite for reading this thesis, as we will only mention it to sketch the context of our results rather than to use it as a tool. In both cases of FO and MSO, we assume that the syntax and semantics are defined in a standard way. Given a formula φ with free variables $x_1, ..., x_k$, and a model $\mathcal{M}$ with points $\mathsf{m}_1, ..., \mathsf{m}_k \in M$, we denote by $\varphi(\mathsf{m}_1, ..., \mathsf{m}_k)$ the formula under interpretation mapping each x_i to m_i. A formula with no free variables is called a *sentence*.

2.4 Fixpoint Extension of ML

We recall the well-known extension of modal logic with fixpoints and the corresponding notions of games and automata (the lecture notes [37] are a good introduction). Apart from setting notation and recalling classical facts, we also include some of the proofs. This will not only make our presentation more self-contained, but also describe the classical setting from a perspective that will facilitate understanding of its extended versions later. It will also make the comparison between the standard and the new setting more transparent.

2.4.1 Fixpoint Modal Logic

Consider a modal formula φ using an atomic proposition τ and a model $\mathcal{M}$ with an interpretation for all the accessibility relations and propositions except for $\tau^{\mathcal{M}}$. One can think of the semantics of φ as a map F that takes the missing interpretation of τ and returns the interpretation of φ, thus transforming subsets of M into subsets of M. If it so happens that τ appears only positively in φ (i.e. it is never negated) then such a map is monotone, meaning that the more points satisfy τ, the more points satisfy φ. Since the powerset $\mathcal{P}(M)$ of M ordered by set inclusion is a complete lattice, by the Knaster-Tarski Theorem 2.1.1 monotonicity of F implies that it has both the least and the greatest fixpoints $\mathsf{LFP}.F = F^{\infty}_{\mu}$ and $\mathsf{GFP}.F = F^{\infty}_{\nu}$. The modal fixpoint logic, called μ-calculus and abbreviated μ-ML, is then obtained from the ordinary modal logic by extending it with operators interpreted as such fixpoints.

Note that in our example the atomic proposition τ was treated as a variable with interpretation ranging over subsets of the model. In order to make the distinction between such variables (thought of as intermediate objects) and actual atomic propositions (that are a fixed component of the model under consideration) we will distinguish variables and atomic propositions on a syntactic level. Thus, from now on we assume an infinite set Var of variables whose elements will be typically denoted x, y, z. The syntax of μ-ML is then given by the grammar:

$$\varphi ::= \top \mid \bot \mid \tau \mid \neg\tau \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \langle \mathsf{a} \rangle \varphi \mid [\mathsf{a}]\varphi \mid x \mid \mu x.\varphi \mid \nu x.\varphi$$

with $\tau \in \mathsf{Prop}$, $\mathsf{a} \in \mathsf{Act}$ and $x \in \mathsf{Var}$. In the last two clauses, the operators μ and ν are said to *bind* the variable x. We call an occurrence of x in φ *bound* if it has a superformula $\psi \in \mathsf{SubFor}(\varphi)$ labelled with an operator binding x. The subformula binding this occurrence is the least such ψ. An occurrence of x is *free* if it is not bound. One can always rename variables that occur bound to fresh ones (e.g. rewrite $x \wedge \mu x.x \wedge \nu x.x$ to $x \wedge \mu y.y \wedge \nu z.z$) because such variable replacement, called *alpha-conversion*, does not change the meaning of formulae nor any of its properties that we investigate. Thus, for technical convenience unless stated otherwise we assume that occurrences of x in φ are either all bound in precisely one place or all free. In particular, no variable is both bound and free in one formula. A *sentence* is then a formula with no free variables. In statements that apply both to least and greatest fixpoints, we will sometimes use η to denote either μ or ν.

For the semantics of μ-ML in a model $\mathcal{M}$ we assume a *valuation* $\mathsf{val} : \mathsf{Var} \to \mathcal{P}(M)$ that interprets all the variables. The semantics $[\![\varphi]\!]^{\mathsf{val}} \subseteq M$ of a formula φ under val is then defined inductively:

$$\begin{aligned}
[\![\top]\!]^{\mathsf{val}} &= M \quad \text{and} \quad [\![\bot]\!]^{\mathsf{val}} = \emptyset \\
[\![\tau]\!]^{\mathsf{val}} &= \tau^{\mathcal{M}} \quad \text{and} \quad [\![\neg\tau]\!]^{\mathsf{val}} = M - \tau^{\mathcal{M}} \\
[\![\varphi_1 \vee \varphi_2]\!]^{\mathsf{val}} &= [\![\varphi_1]\!]^{\mathsf{val}} \cup [\![\varphi_2]\!]^{\mathsf{val}} \quad \text{and} \quad [\![\varphi_1 \wedge \varphi_2]\!]^{\mathsf{val}} = [\![\varphi_1]\!]^{\mathsf{val}} \cap [\![\varphi_2]\!]^{\mathsf{val}} \\
[\![\langle \mathsf{a} \rangle \varphi]\!]^{\mathsf{val}} &= \{\mathsf{m} \in M \mid \exists_{\mathsf{m} \xrightarrow{\mathsf{a}} \mathsf{n}}\ \mathsf{n} \in [\![\varphi]\!]^{\mathsf{val}}\} \quad \text{and} \quad [\![[\mathsf{a}]\varphi]\!]^{\mathsf{val}} = \{\mathsf{m} \in M \mid \forall_{\mathsf{m} \xrightarrow{\mathsf{a}} \mathsf{n}}\ \mathsf{n} \in [\![\varphi]\!]^{\mathsf{val}}\} \\
[\![x]\!]^{\mathsf{val}} &= \mathsf{val}(x) \\
[\![\mu x.\varphi]\!]^{\mathsf{val}} &= F^{\infty}_{\mu} \quad \text{and} \quad [\![\nu x.\varphi]\!]^{\mathsf{val}} = F^{\infty}_{\nu}
\end{aligned}$$

where in the last clause $F(H) = [\![\varphi]\!]^{\mathsf{val}[x \mapsto H]}$. We will skip the index val if it is immaterial or clear from the context. The *closure ordinal* of φ is the least $\alpha \in \mathsf{Ord}_{\infty}$ such that for all models and valuations the induced map F reaches its fixpoints in α-many steps, meaning that $F^{\alpha}_{\mu} = F^{\infty}_{\mu}$ and $F^{\alpha}_{\nu} = F^{\infty}_{\nu}$.

Example 2.4.1. Consider the monomodal fixpoint formula $\varphi = \mu x.\Box x$. Given a model $\mathcal{M}$, its semantics $[\![\varphi]\!]$ equals to the least fixpoint of the map $F(H) = [\![\Box x]\!]^{\mathsf{val}[x \mapsto H]} = \{\mathsf{m} \in M \mid \forall_{\mathsf{m} \to \mathsf{n}}.\ \mathsf{n} \in H\}$. This F is the operation $F_{\Box}$ for the structure $(M, \to^{\mathcal{M}})$ and hence by (2.1) we get that φ is true at a point iff no infinite path starts there.

Negation Note that, similarly as with ML, the syntax of μ-ML does not include negation, but its sentences are semantically closed under negation in the following sense. For every formula φ there is a dual formula $\widetilde{\varphi}$ such that:

$$[\![\widetilde{\varphi}]\!]^{\mathsf{val}} = M - [\![\varphi]\!]^{\widetilde{\mathsf{val}}}$$

for every model $\mathcal{M}$ and valuations val and $\widetilde{\mathsf{val}}$ such that $\widetilde{\mathsf{val}}(x) = M - \mathsf{val}(x)$ for all $x \in \mathsf{Var}$. If φ is a sentence, then both val and $\widetilde{\mathsf{val}}$ are irrelevant and so the semantics of $\widetilde{\varphi}$ is the complement of the semantics of φ. Such $\widetilde{\varphi}$ is constructed by induction on the complexity of φ with all the boolean and modal cases defined the same as for plain ML and:

$$\widetilde{x} = x$$
$$\widetilde{\mu x.\varphi} = \nu x.\widetilde{\varphi}, \qquad \text{and} \qquad \widetilde{\nu x.\varphi} = \mu x.\widetilde{\varphi}$$

Correctness of such rewriting follows from the Knaster-Tarski characterization of fixpoints from Theorem 2.1.1 and the de Morgan laws:

$$\neg \bigvee A = \bigwedge\{\neg a \mid a \in A\} \quad \text{and} \quad \neg \bigwedge A = \bigvee\{\neg a \mid a \in A\}$$

which are satisfied by the powerset lattice $\mathcal{P}(M)$ with $\neg$ interpreted as the complement.

Example 2.4.2. Consider the formula $\varphi = \mu x.\Box x$ from Example 2.4.1 which defines well-foundedness. According to the definition, the dual formula $\widetilde{\varphi}$ is:

$$\widetilde{\mu x.\Box x} = \nu x.\widetilde{\Box x} = \nu x.\Diamond\widetilde{x} = \nu x.\Diamond x.$$

The reader may check that indeed $\nu x.\Diamond x$ is true at a point iff there is an infinite path starting there.

Vectorial μ-calculus. A syntactically richer version of the modal μ-calculus admits mutual fixpoint definitions of multiple properties, in formulae such as:

$$\mu_1(x_1, x_2).(\varphi_1, \varphi_2),$$

where variables x_1 and x_2 may occur both in φ_1 and φ_2. Given a valuation val as before, this formula is interpreted as the least fixpoint of the monotone function $(H_1, H_2) \mapsto ([\![\varphi_1]\!]^{\mathsf{val}[x_i \mapsto H_i]}, [\![\varphi_1]\!]^{\mathsf{val}[x_i \mapsto H_i]})$ on the complete lattice $\mathcal{P}(M) \times \mathcal{P}(M)$; the resulting pair of sets is then projected to the first component as dictated by the subscript in μ_1. Tuples of any size are allowed. This *vectorial* calculus is expressively equivalent to the scalar version described before, thanks to the so-called *Bekić principle* ([1, Bisection Lemma]) which says that the equality:

$$\mu \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} . \begin{pmatrix} f_1(x_1, x_2) \\ f_2(x_1, x_2) \end{pmatrix} = \begin{pmatrix} \mu x_1.f_1(x_1,\ \mu x_2.f_2(x_1, x_2)) \\ \mu x_2.f_2(\mu x_1.f_1(x_1, x_2),\ x_2) \end{pmatrix} \tag{2.9}$$

holds for every pair of monotone operations $f_i : A_1 \times A_2 \to A_i$ on complete lattices A_1, A_2, and similarly for the greatest fixpoint operator ν in place of μ.

2.4.2 Simple Parity Games

The μ-calculus has a straightforward inductively-defined semantics, but it is often useful to consider an alternative (but equivalent) semantics based on parity games. Among other advantages, the game semantics provides more efficient algorithms for model checking of μ-calculus formulas than an inductive computation of fixpoints [11].

We have already seen in Subsection 2.2.3 how fixpoint games characterize fixpoints of monotone operations. In this case the winning condition was as simple as it can be: ∃ve looses all the infinite plays in the least fixpoint game $\mathcal{G}_\mu$ and wins them in the greatest fixpoint game $\mathcal{G}_\nu$. In the context of μ-ML the fixpoint operators μ and ν can be nested. Because of this, in order to capture the meaning of formulae with a game we need the parity winning condition that reflects such nesting.

In the basic classical setting that we recall now, configurations of a game do not possess any relevant internal structure and hence can be identified with positions. Recall that such games are called *simple*.

Definition 2.4.3. Configurations of a *simple parity game* are identified with positions from a fixed arena:

$$\mathsf{Conf} = V = V_\exists \sqcup V_\forall$$

divided between the players. The legal moves are given by an *edge relation*:

$$\mathsf{Mov} = E \subseteq V \times V$$

and the winning condition by a *rank function* $\mathsf{rank} : V \to \mathcal{R}$ that maps each position to an element of a fixed finite linear order $\mathcal{R} = \mathcal{R}_\exists \sqcup \mathcal{R}_\forall$ divided between ∃ve and ∀dam.

2.4.3 Parity Automata

The notion of parity game gives rise to parity *automata*, a stepping stone between logic and games. The connection is threefold.

First, every μ-ML-formula φ defines an automaton $\mathcal{A}_\varphi$ that recognizes the same language. In fact, $\mathcal{A}_\varphi$ arises from φ in a very direct way, as it suffices to define an appropriate automaton structure on the set $\mathsf{SubFor}(\varphi)$ of subformulae of φ. The construction is designed so that the semantic game $\mathcal{G}(\mathcal{A}_\varphi)$ extends the game semantics for plain ML with extra rules for the new operators. Since the translation allows us to view a formula as an automaton rather than to construct an abstract device, it is often referred to as the *game semantics* for μ-ML.

Second, automata can be seen as a relaxation of the notion of a formula seen as a game-inducing object, where only the elements that are crucial from that point of view are kept. It then turns out that such generalization, although often convenient, does not increase the expressive power of the model: for every abstract automaton $\mathcal{A}$ there is a formula $\varphi_{\mathcal{A}}$ defining the same language.

Third, an arbitrary game $\mathcal{G} = (V, E, \mathsf{rank})$ can be seen as a modal model. Then, the logic (or equivalently: automata) can solve the game in the sense that the set of wining positions of $\mathcal{G}$ is definable by a formula.

Since formulae can have free variables, for technical reasons we will also consider automata with free variables. These variables resemble terminal states in that they can be targets of transitions, but no transitions originate in them, and whether they accept or not depends on an external valuation. Moreover, we want our automata to be able to test whether a given atomic proposition τ is true or not. Thus, we also include τ and $\neg\tau$, indicating such test, as legal targets of transitions.

A parity automaton consists of:

- a finite set of states $Q = Q_\exists \sqcup Q_\forall$ divided between two players;
- an initial state $q_I \in Q$;
- a transition function:

$$\delta : Q \to \mathcal{P}(Q \sqcup \mathsf{Lit} \sqcup \mathsf{Var}) \sqcup (\mathsf{Act} \times Q)$$

  where the first part is called *ε-transitions* and the second one *modal transitions*;
- a function assigning ranks to states $\mathsf{rank} : Q \to \mathcal{R}$.

The language of an automaton is a subclass of all pointed models, defined in terms of parity games. Fix a model $\mathcal{M}$. Given an automaton $\mathcal{A} = (Q, q_I, \delta, \mathsf{rank})$ and a valuation $\mathsf{val} : \mathsf{Var} \to \mathcal{P}(M)$, we define the semantic game $\mathcal{G}^{\mathsf{val}}(\mathcal{A})$ to be the parity game (V, E, rank') where positions are of the form:

$$V = M \times (Q \sqcup \mathsf{Lit} \sqcup \mathsf{Var})$$

and the edge relation E is defined as follows. In a position (m, q) for $q \in Q$:

- if $\delta(q) \subseteq Q \sqcup \mathsf{Lit} \sqcup \mathsf{Var}$, outgoing edges (called ε-edges, or ε-moves) are:

$$\{((\mathsf{m}, q), (\mathsf{m}, z)) \mid z \in \delta(q)\},$$

- if $\delta(q) = (\mathsf{a}, p)$, outgoing edges (modal edges, modal moves) are:
$$\{((\mathsf{m}, q), (\mathsf{n}, p)) \mid \mathsf{m} \xrightarrow{\mathsf{a}} \mathsf{n}\}.$$

There are no outgoing edges from positions $(\mathsf{m}, x), (\mathsf{m}, \tau)$ nor $(\mathsf{m}, \neg\tau)$ for $x \in \mathsf{Var}$ and $\tau \in \mathsf{Prop}$, which means that the owner of these positions looses immediately.

For $q \in Q$, the owner of the position (m, q) is the owner of the state q, and $\mathsf{rank}'(\mathsf{m}, q) = \mathsf{rank}(q)$. For $x \in \mathsf{Var}$, the position (m, x) belongs to $\forall$dam if $\mathsf{m} \in \mathsf{val}(x)$ and to $\exists$ve otherwise. Similarly, (m, τ) (or $(\mathsf{m}, \neg\tau)$) belongs to $\exists$ve iff $\mathsf{m} \notin \tau^{\mathcal{M}}$ (or $\mathsf{m} \in \tau^{\mathcal{M}}$, respectively). The rank' of (m, x), (m, τ) and $(\mathsf{m}, \neg\tau)$ can be set arbitrarily, as it does not affect the outcome of the game. The semantics $[\![\mathcal{A}]\!]^{\mathsf{val}} \subseteq M$ of an automaton $\mathcal{A}$ is the set of all points $\mathsf{m} \in M$ for which the position (m, q_I) in the game $\mathcal{G}^{\mathsf{val}}(\mathcal{A})$ is winning for $\exists$ve, in which case we say that $\mathcal{A}$ *accepts* $\mathcal{M}, \mathsf{m}$. The language of $\mathcal{A}$ is the class of all pointed models it accepts.

Remark 2.4.4. It should be emphasized that, unless stated otherwise, by automata we always mean *alternating* ones. That is, in the semantic game we allow both players to make decisions. Alternating automata can be contrasted with its subclasses of *nondeterministic* and *deterministic* ones, where all the choices are made by $\exists$ve, or there are no nontrivial choices at all, respectively. In some cases these models are equivalent, meaning that given an alternating automaton one can compute a nondeterministic, or even deterministic automaton recognizing the same language. However, in general the alternating model is strictly more expressive.

We now inspect the translations between automata and logic. Both are classical and can be found e.g. in [37]. However, since the standard constructions and proofs form a basis for more complex generalizations later, we recall them now, both to put the new results into context and to set notation.

2.4.4 From Formulae to Automata

Every fixpoint formula $\varphi \in \mu\text{-}\mathsf{ML}$ gives rise to a parity automaton $\mathcal{A}_\varphi$ such that $[\![\varphi]\!]^{\mathsf{val}} = [\![\mathcal{A}_\varphi]\!]^{\mathsf{val}}$ for every model $\mathcal{M}$ and valuation val. Specifically, given a formula φ (with some free variables), we define an automaton $\mathcal{A}_\varphi = (Q, q_I, \delta, \mathsf{rank})$ (over the same free variables) as follows:

- Q is the set of all subformulae other than the literals and free variables of φ:
$$Q = \mathsf{SubFor}(\varphi) - (\mathsf{Lit} \sqcup \mathsf{FreeVar}(\varphi))$$
Ownership of a state in Q depends on the topmost connective, with $\exists$ve owning $\vee$ and $\langle \mathsf{a} \rangle$ and $\forall$dam owning $\wedge$ and $[\mathsf{a}]$; ownership of fixpoint subformulae can be set arbitrarily as it will not matter;

- $q_I = \varphi$;
- the transition function is defined by cases:
  - $\delta(\theta_1 \vee \theta_2) = \delta(\theta_1 \wedge \theta_2) = \{\theta_1, \theta_2\}$,
  - $\delta(\langle \mathsf{a} \rangle \theta) = \delta([\mathsf{a}]\theta) = (\mathsf{a}, \theta)$,
  - $\delta(\eta x.\theta) = \{\theta\}$ (for $\eta = \mu$ or $\eta = \nu$),
  - $\delta(x) = \{\theta\}$, where $\eta x.\theta$ is the (unique) subformula of φ binding x.
- For the ranking function rank, take the lowest rank 0 (its ownership does not matter) and assign it to all subformulae of φ except for immediate subformulae of fixpoint operators (i.e. θ for some $\eta x.\theta$). To those, assign ranks in such a way that subformulae have strictly smaller ranks than their superformulae, and for every subformula $\eta x.\theta$ the rank of θ belongs to $\exists$ve if $\eta = \mu$ and to $\forall$dam if $\eta = \nu$.

We denote $\mathcal{G}^{\mathsf{val}}(\varphi) = \mathcal{G}^{\mathsf{val}}(\mathcal{A}_\varphi)$.

Remark 2.4.5. Recall that a subformula of φ is a node in its syntactic tree. Hence, in the definition above clauses like $\delta(\theta_1 \vee \theta_2) = \{\theta_1, \theta_2\}$ should be read as: the value of δ on a node labelled by disjunction $\vee$ equals the set of immediate children of that node. The only possibly confusing case is the last one for a bound variable x. It states that the target of a transition originating in an occurrence of x (i.e. a node labelled with x) is the immediate subformula of the formula binding that occurrence (i.e. the child of the closest ancestor of this occurrence of x labelled with ηx). Under our usual assumption all the occurrences of each bound variable are bound in the same place. In that case the definition is uniform for all occurrences of x. This special uniform case is already general in the sense that every formula can be alpha-converted to equivalent one satisfying the assumption and alpha-conversion does not change the induced automata (up to isomorphism). Nonetheless, it will be convenient to occasionally consider automata arising from formulae with the same variable bound in multiple places. Formulae with the same variable occurring both free and bound will not be considered.

Remark 2.4.6. Every automaton $\mathcal{A}$, when paired with a model $\mathcal{M}$ and valuation val, induces the semantic game $\mathcal{G}^{\mathsf{val}}(\mathcal{A})$. The above translation $\varphi \mapsto \mathcal{A}_\varphi$ is crafted so that the compositional definition of semantic games for purely modal formulae (i.e. with no fixpoint operators) defined in Subsection 2.3 is a special case of the one for μ-ML. To that end, the game for ML is extended with rules determining what happens in configurations having $x \in \mathsf{Var}$ or $\eta x.\theta$ on the formula coordinate. Since the new rules result in a possibility of infinite plays, we need to determine who wins in such cases. In every infinite play we look at the outermost operator η that was seen infinitely often and classify the play as won by $\exists$ve iff $\eta = \nu$.

Theorem 2.4.7 (μ-ML Adequacy). *For every model $\mathcal{M}$ and valuation* val, $[\![\varphi]\!]^{\mathsf{val}} = [\![\mathcal{A}_\varphi]\!]^{\mathsf{val}}$.

Proof. Unfolding the definition of $[\![\mathcal{A}_\varphi]\!]^{\mathsf{val}}$ we prove that:

$$\mathsf{m} \in [\![\varphi]\!]^{\mathsf{val}} \iff \exists\text{ve wins } \mathcal{G}^{\mathsf{val}}(\varphi) \text{ from } (\mathsf{m}, \varphi) \tag{2.10}$$

for every $\mathsf{m} \in M$ and valuation val. The proof proceeds by induction on the complexity of φ. All the cases except for $\varphi = \mu x.\psi$ and $\varphi = \nu x.\psi$ are easy and the proof does not differ from the one for plain ML expressed by Theorem 2.3.1. Let us focus on the only remaining case with $\varphi = \mu x.\psi$ (the case with $\varphi = \nu x.\psi$ is symmetric). Since by definition $[\![\varphi]\!]^{\mathsf{val}} = \mathsf{LFP}.F$ for $F(H) = [\![\psi]\!]^{\mathsf{val}[x \mapsto H]}$, Theorem 2.2.8 gives us:

$$\mathsf{m} \in [\![\varphi]\!]^{\mathsf{val}} \iff \mathsf{m} \in \mathsf{LFP}.F \iff \exists\text{ve wins } \mathcal{G}_\mu(F) \text{ from } \mathsf{m}.$$

Hence, in order to prove (2.10) it suffices if we show the exit-equivalence:

$$\mathcal{G}_\mu(F), \mathsf{m} \bowtie \mathcal{G}^{\mathsf{val}}(\varphi), (\mathsf{m}, \varphi)$$

for all $\mathsf{m} \in M$. We would like to decompose $\mathcal{G}_\mu(F)$ and $\mathcal{G}^{\mathsf{val}}(\varphi)$ into phases ending in M and $M \times \{\psi\}$, respectively, and apply the Decomposition Lemma 2.2.5 (with equivalence of such phases following from the induction hypothesis for ψ). To that end we view $\mathcal{G}_\mu(F)$ as a reachability game. Denote:

$$\mathcal{G} = \mathcal{G}_\mu(F) \qquad \text{and} \qquad \mathcal{G}' = \mathcal{G}^{\mathsf{val}}(\varphi).$$

Since $\varphi = \mu x.\psi$, it follows that $\mathcal{G}'$ moves deterministically from position (m, φ) to (m, ψ) and therefore it suffices to prove:

$$\mathcal{G}, \mathsf{m} \bowtie \mathcal{G}', (\mathsf{m}, \psi). \tag{2.11}$$

Towards the use of the lemma, define:

$$\mathcal{S}_+ = M \qquad \text{and} \qquad \mathcal{S}'_+ = M \times \{\psi\}$$

and:

$$S_+ = \{(\mathsf{n}, (\mathsf{n}, \psi)) \mid \mathsf{n} \in M\} \subseteq \mathcal{S}_+ \times \mathcal{S}'_+$$

linking every n to (n, ψ). $\mathcal{S}_+$ and $\mathcal{S}'_+$ only contain configurations with the highest ranks of $\mathcal{G}$ and $\mathcal{G}'$, respectively. Hence, substituting $S = \emptyset$, we get (2.11) from the Decomposition Lemma 2.2.5 provided that we prove:

$$\mathcal{G}, \mathsf{n} \bowtie_{S_+} \mathcal{G}', (\mathsf{n}, \psi)$$

for all $\mathsf{n} \in M$. For that, we show that for every $H \subseteq M$ the following are equivalent:

1. $\exists$ve has a non-loosing strategy for $\mathcal{G}|\mathcal{S}_+, \mathsf{n}$ with exit configurations H,

2. $\mathsf{n} \in F(H)$,

3. $\exists$ve has a winning strategy for $\mathcal{G}^{\mathsf{val}[x\mapsto H]}(\psi), (\mathsf{n}, \psi)$,

4. $\exists$ve has a non-loosing strategy for $\mathcal{G}'|\mathcal{S}'_+, (\mathsf{n}, \psi)$ with exit configurations included in $H \times \{\psi\}$.

This proves exit-equivalence: the implication (1) $\Longrightarrow$ (4) gives an immediate translation of strategies σ in $\mathcal{G}|\mathcal{S}_+, \mathsf{n}$ to σ' in $\mathcal{G}'|\mathcal{S}'_+, (\mathsf{n}, \psi)$ with $\mathsf{exit}(\sigma') \subseteq \mathsf{exit}(\sigma)\times\{\psi\} = S_+[\mathsf{exit}(\sigma)]$. For the opposite direction, take σ' in $\mathcal{G}'|\mathcal{S}'_+, (\mathsf{n}, \psi)$. Consider $H = \pi_1(\mathsf{exit}(\sigma'))$ so that $\mathsf{exit}(\sigma') \subseteq H \times \{\psi\}$. We get from (4) $\Longrightarrow$ (1) a strategy σ for $\mathcal{G}|\mathcal{S}_+, \mathsf{n}$ with $\mathsf{exit}(\sigma) = H = S_+^{-1}[\mathsf{exit}(\sigma')]$. We complete the proof of the Theorem 2.4.7 by proving the equivalence of all four items (1), (2), (3) and (4).

For (1) $\Longleftrightarrow$ (2) observe that in $\mathcal{G}|\mathcal{S}_+, \mathsf{n}$ $\exists$ve picks a subset $H \subseteq M$ such that $\mathsf{n} \in F(H)$, then $\forall$dam chooses its element $\mathsf{n}' \in H$ and the game stops. Thus, strategies for $\exists$ve can be identified with such subsets of M (each such H viewed as a strategy is non-loosing and has H itself as the set of its exit configurations).

Since by definition $F(H) = [\![\psi]\!]^{\mathsf{val}[x\mapsto H]}$, (2) $\Longleftrightarrow$ (3) follows from the induction hypothesis (2.10) applied to ψ.

For (3) $\Longleftrightarrow$ (4) note that the games $\mathcal{G}'|\mathcal{S}'_+, (\mathsf{n}, \psi)$ and $\mathcal{G}^{\mathsf{val}[x\mapsto H]}(\psi), (\mathsf{n}, \psi)$ do not differ until a move to a position of shape (n', x). If it happens, $\mathcal{G}'|\mathcal{S}'_+$ moves deterministically to $(\mathsf{n}', \psi) \in \mathcal{S}'_+$ and stops with no winner whereas $\mathcal{G}^{\mathsf{val}[x\mapsto H]}(\psi)$ ends and $\exists$ve wins iff $\mathsf{n}' \in H$. It follows that winning strategies for $\exists$ve in $\mathcal{G}^{\mathsf{val}[x\mapsto H]}(\psi), (\mathsf{n}, \psi)$ are the same as her non-loosing strategies in $\mathcal{G}'|\mathcal{S}'_+, (\mathsf{n}, \psi)$ whose exit configurations are included in $H \times \{\psi\}$. $\square$

2.4.5 From Automata to Formulae

Theorem 2.4.8. *For every parity automaton $\mathcal{A}$ there exists a formula $\varphi_{\mathcal{A}}$ of μ-ML such that $[\![\mathcal{A}]\!]^{\mathsf{val}} = [\![\varphi_{\mathcal{A}}]\!]^{\mathsf{val}}$ for every model $\mathcal{M}$ and valuation* val.

Proof. Fix an automaton $\mathcal{A} = (Q, q_I, \delta, \mathsf{rank})$. For clarity of presentation we only consider the case when $\mathcal{A}$ has no free variables, the general case requires no new ideas. Without losing generality assume that the highest rank r_{max} is not assigned to any state and every other rank is assigned to precisely one state: unused ranks can be removed and every rank r assigned to multiple states $q_1, ..., q_k$ can be replaced with a linearly ordered sequence of its copies $r_1 \preceq ... \preceq r_k$, one for each state, as this does not change the winner of any play. Take variables $\mathsf{Var}_Q = \{x_q \mid q \in Q\}$ with distinct $x_q \in \mathsf{Var}_Q$ for every $q \in Q$ and denote:

$$Q_{r\leq} = \{q \in Q \mid r \leq \mathsf{rank}(q)\} \qquad \text{and} \qquad \mathsf{Var}_{r\leq} = \{x_q \mid q \in Q_{r\leq}\}.$$

We construct, by induction on $r \in \mathcal{R}$, a formula $\psi_{r,q}$ over Var_Q with all free variables in $\mathsf{Var}_{r\leq}$ and all bound variables outside of $\mathsf{Var}_{r\leq}$. The goal of our construction is that for every point $\mathsf{m} \in \mathcal{M}$:

$$\mathcal{G}(\mathcal{A}), (\mathsf{m}, q) \bowtie_{S_{r\leq}} \mathcal{G}(\psi_{r,q}), (\mathsf{m}, \psi_{r,q}) \tag{2.12}$$

where:

$$S_{r\leq} = \{((\mathsf{n}, p), (\mathsf{n}, x_p)) \mid p \in Q_{r\leq}\}$$

with domain and codomain:

$$\mathcal{S}_{r\leq} = M \times Q_{r\leq} \qquad \text{and} \qquad \mathcal{S}'_{r\leq} = M \times \mathsf{Var}_{r\leq}.$$

That is, $\psi_{r,q}$ is designed so that the games $\mathcal{G}(\mathcal{A}), (\mathsf{m}, q)$ and $\mathcal{G}(\psi_{r,q}), (\mathsf{m}, \psi_{r,q})$ are the same *until* a move to a state p or a variable x_p, respectively, such that the rank of p is at least r. Hence, (2.12) for the highest rank r_{max} implies the theorem: since no state in $\mathcal{A}$ has rank r_{max}, the relation $S_{r_{\mathsf{max}}\leq} = \emptyset$ is empty and (2.12) means that the games $\mathcal{G}(\mathcal{A}), (\mathsf{m}, q_I)$ and $\mathcal{G}(\psi_{r_{\mathsf{max}},q_I}), (\mathsf{m}, \psi_{r_{\mathsf{max}},q_I})$ are equivalent.

Note that although formally $\psi_{r,q}$ may contain free variables, the partial game $\mathcal{G}^{\mathsf{val}}(\psi_{r,q})|\mathcal{S}'_{r\leq}, (\mathsf{m}, \psi_{r,q})$ always stops when any such variable is reached. Thus, we ignore the valuation val as irrelevant and write $\mathcal{G}(\psi_{r,q})$. Moreover, the reader should be warned that the constructed formulae do not have to satisfy our usual assumption about bound variables: a variable x_p can be bound by multiple operators in $\psi_{r,q}$. One could avoid that as follows: whenever we build a formula θ using already constructed $\theta_1, ..., \theta_l$, first replace all the bound variables in each used copy of each θ_k with fresh ones. Nonetheless, to avoid clumsy notation we refrain ourselves from doing that.

The Base Case. Consider the lowest rank 0. The set $\mathcal{S}_{0\leq} = M \times Q$ contains all the positions of $\mathcal{G}(\mathcal{A})$ except for the ones from $M \times (\mathsf{Lit} \sqcup \mathsf{Var})$. This means that after the first move $\mathcal{G}(\mathcal{A})|\mathcal{S}_{0\leq}$ either stops immediately or reaches such a terminal position. Thus for the base case of (2.12) it is enough to put:

- if $\delta(q) = (\mathsf{a}, p)$:

$$\psi_{0,q} = \begin{cases} \langle \mathsf{a} \rangle p & \text{if } q \text{ belongs to } \exists\text{ve} \\ [\mathsf{a}]p & \text{if } q \text{ belongs to } \forall\text{dam} \end{cases}$$

- if $\delta(q) \subseteq Q \sqcup \mathsf{Lit} \sqcup \mathsf{Var}$:

$$\psi_{0,q} = \begin{cases} \bigvee \delta(q) & \text{if } q \text{ belongs to } \exists\text{ve} \\ \bigwedge \delta(q) & \text{if } q \text{ belongs to } \forall\text{dam}. \end{cases}$$

The Inductive Step. For the inductive step, assume that for some rank r and each $q \in Q$ we have a formula $\psi_{r,q}$ satisfying (2.12). Denoting the next rank by $r + 1$ we construct $\psi_{r+1,q}$ so that:

$$\mathcal{G}(\mathcal{A}), (\mathsf{m}, q) \bowtie_{S_{r+1\leq}} \mathcal{G}(\psi_{r+1,q}), (\mathsf{m}, \psi_{r+1,q}) \tag{2.13}$$

for all $\mathsf{m} \in \mathcal{M}$. Let p be the unique state in Q that has rank r. In the inductive proof we use the decomposition:

$$S_{r\leq} = S_r \sqcup S_{r+1\leq} \qquad \text{with} \qquad S_r = \{((\mathsf{n}, p), (\mathsf{n}, x_p)) \mid \mathsf{n} \in M\}.$$

To simplify notation denote:

$$S = S_{r+1\leq} \subseteq \mathcal{S} \times \mathcal{S}' \quad \text{where} \quad \mathcal{S} = \mathcal{S}_{r+1\leq} \quad \text{and} \quad \mathcal{S}' = \mathcal{S}'_{r+1\leq}.$$

The Case with $q = p$. We start the construction with the case $q = p$ and put:

$$\psi_{r+1,p} = \eta x_p.\psi_{r,p}$$

where $\eta = \mu$ if r belongs to $\exists$ve and $\eta = \nu$ if r belongs to $\forall$dam. The idea is that the formula $\psi_{r+1,p}$ induces the same game as $\psi_{r,p}$ except that upon a visit to any (n, x_p) the play continues instead of stopping. Since by the induction hypothesis the play leading from $(\mathsf{m}, \psi_{r,p})$ to (n, x_p) in $\mathcal{G}(\psi_{r,p})$ corresponds to a play leading from (m, p) to (n, p) in $\mathcal{G}(\mathcal{A})$, we would like the new game to continue from $(\mathsf{n}, \psi_{r,p})$. The above definition does exactly that: the fixpoint operator η bounding x_p adds a deterministic transition from (n, x_p) to $(\mathsf{n}, \psi_{r,p})$. The choice of μ or ν guarantees that in case of infinitely many such unfoldings the winner in $\mathcal{G}(\psi_{r+1,p})$ is the same as in $\mathcal{G}(\mathcal{A})$.

Formally, we claim that:

$$\mathcal{G}(\mathcal{A}), (\mathsf{m}, p) \bowtie_S \mathcal{G}(\psi_{r+1,p}), (\mathsf{m}, \psi_{r,p}) \tag{2.14}$$

for all $\mathsf{m} \in M$. Since $\psi_{r+1,p} = \eta x_p.\psi_{r,p}$, the game $\mathcal{G}(\psi_{r+1,p})$ moves deterministically from $(\mathsf{m}, \psi_{r+1,p})$ to $(\mathsf{m}, \psi_{r,p})$ and hence (2.14) implies the induction goal (2.13). Define sets of configurations:

$$\begin{aligned} \mathcal{S}_+ &= M \times \{p\} \\ \mathcal{S}_+^\circ &= M \times \{x_p\} \\ \mathcal{S}'_+ &= M \times \{\psi_{r,p}\} \end{aligned}$$

and relations $R \subseteq \mathcal{S}_+^\circ \times \mathcal{S}'_+$ and $S_+ \subseteq \mathcal{S}_+ \times \mathcal{S}'_+$:

$$\begin{aligned} R &= \{((\mathsf{n}, x_p), (\mathsf{n}, \psi_{r,p})) \mid \mathsf{n} \in M\} \\ S_+ &= \{((\mathsf{n}, p), (\mathsf{n}, \psi_{r,p})) \mid \mathsf{n} \in M\}. \end{aligned}$$

We then have:

$$\begin{aligned} \mathcal{G}(\mathcal{A}), (\mathsf{m}, p) &\bowtie_{S_{r\leq}} \mathcal{G}(\psi_{r,p}), (\mathsf{m}, \psi_{r,p}) \\ &\bowtie_{R \sqcup \mathsf{Id}(\mathcal{S}')} \mathcal{G}(\psi_{r+1,p}), (\mathsf{m}, \psi_{r,p}) \end{aligned}$$

for all $\mathsf{m} \in \mathcal{M}$. The first equivalence is the induction hypothesis (2.12). The second one is true because the partial games $\mathcal{G}(\psi_{r,p})|\mathcal{S}_+^\circ \sqcup \mathcal{S}', (\mathsf{m}, \psi_{r,p})$ and

$\mathcal{G}(\psi_{r+1,p})|\mathcal{S}'_+ \sqcup \mathcal{S}', (\mathsf{m}, \psi_{r,p})$ are isomorphic until a move to some $(\mathsf{n}, x_p) \in \mathcal{S}^\circ_+$ in which case the first one stops and the second one moves deterministically to $(\mathsf{n}, \psi_{r,p}) \in \mathcal{S}'_+$ and stops as well. We have $S_{r\leq} = S_r \sqcup S$, $S_r \circ R = S_+$ and $S \circ \mathsf{Id}(\mathcal{S}') = S$ and therefore:

$$\begin{aligned} S_{r\leq} \circ (R \sqcup \mathsf{Id}(\mathcal{S}')) &= (S_r \sqcup S) \circ (R \sqcup \mathsf{Id}(\mathcal{S}')) \\ &= (S_r \circ R) \sqcup (S \circ \mathsf{Id}(\mathcal{S}')) \\ &= S_+ \sqcup S. \end{aligned}$$

This allows us to compose the equivalences into:

$$\mathcal{G}(\mathcal{A}), (\mathsf{m}, p) \bowtie_{S_+ \sqcup S} \mathcal{G}(\psi_{r+1,p}), (\mathsf{m}, \psi_{r,p}) \tag{2.15}$$

for all $\mathsf{m} \in \mathcal{M}$. Denoting $\mathcal{G} = \mathcal{G}(\mathcal{A})$ and $\mathcal{G}' = \mathcal{G}(\psi_{r+1,p})$, we get that $\mathcal{S}_+$ and $\mathcal{S}'_+$ only contain configurations with the most important rank r in $\mathsf{Conf} - \mathcal{S}$ and r' in $\mathsf{Conf}' - \mathcal{S}'$, respectively. Moreover, by the choice of $\eta \in \{\mu, \nu\}$, r and r' have the same owner in both games. Hence, applying the Decomposition Lemma 2.2.5 we get (2.14) from (2.15).

The Case with $q \neq p$. It remains to construct formulae $\psi_{r+1,q}$ for all other states $q \neq p$. In this case we take $\psi_{r,q}$ and replace all the occurrences of the free variable x_p with $\psi_{r+1,p}$ constructed in the previous case:

$$\psi_{r+1,q} = \psi_{r,q}[x_p \mapsto \psi_{r+1,p}].$$

As in the previous case, we want to get a formula that induces the same game as $\psi_{r,q}$ except that it does not stop at (n, x_p). The above substitution reflects this idea: in the new game every move to (n, x_p) is replaced with a move to $(\mathsf{n}, \psi_{r+1,p})$. From there, the game is the same as $\mathcal{G}(\psi_{r+1,p}), (\mathsf{n}, \psi_{r+1,p})$ and its correctness is already covered in the previous case.

Let us make the above idea more formal. First observe that $\psi_{r+1,q}$ may contain different copies $\theta_1, ..., \theta_l$ of $\psi_{r+1,p}$ as subformulae. Therefore, we denote subformulae of the k-th copy θ_k with a superscript k, e.g. $\theta_k = \psi^k_{r+1,p}$.

This way, we define relations $S_+ \subseteq \mathcal{S}_+ \times \mathcal{S}'_+$ and $R \subseteq \mathcal{S}^\circ_+ \times \mathcal{S}'_+$ almost the same as in the previous case except that the relations involve all the copies $\psi^1_{r,p}, ..., \psi^l_{r,p}$ of $\psi_{r,p}$ instead of just $\psi_{r,p}$. That is:

$$\begin{aligned} \mathcal{S}_+ &= M \times \{p\} \\ \mathcal{S}^\circ_+ &= M \times \{x_p\} \\ \mathcal{S}'_+ &= M \times \{\psi^k_{r,p} \mid k \leq l\} \end{aligned}$$

and:

$$\begin{aligned} R &= \{((\mathsf{n}, x_p), (\mathsf{n}, \psi^k_{r,p})) \mid k \leq l, \mathsf{n} \in M\} \\ S_+ &= \{((\mathsf{n}, p), (\mathsf{n}, \psi^k_{r,p})) \mid k \leq l, \mathsf{n} \in M\}. \end{aligned}$$

Using the above relations we obtain the induction goal (2.13) from Proposition 2.2.2 substituting $\mathcal{G} = \mathcal{G}(\mathcal{A})$ and $\mathcal{G}' = \mathcal{G}(\psi_{r+1,q})$. The proposition assumes two premises:

$$\mathcal{G}(\mathcal{A}), (\mathsf{m}, q) \bowtie_{S_+ \sqcup S} \mathcal{G}(\psi_{r+1,q}), (\mathsf{m}, \psi_{r+1,q}) \tag{2.16}$$

and:

$$\mathcal{G}(\mathcal{A}), (\mathsf{m}, p) \bowtie_{S} \mathcal{G}(\psi_{r+1,q}), (\mathsf{m}, \psi^k_{r,p}) \tag{2.17}$$

for all $\mathsf{n} \in M$ and $k \leq l$. The first premise (2.16) follows from the composition:

$$\begin{aligned}\mathcal{G}(\mathcal{A}), (\mathsf{m}, q) &\bowtie_{S_{r\leq}} \mathcal{G}(\psi_{r,q}), (\mathsf{m}, \psi_{r,q}) \\ &\bowtie_{R \sqcup \mathsf{Id}(\mathcal{S}')} \mathcal{G}(\psi_{r+1,q}), (\mathsf{m}, \psi_{r+1,q}).\end{aligned}$$

as in the previous case. The first line is the induction hypothesis (2.12). The second one follows from the definition of $\psi_{r+1,q} = \psi_{r,q}[x_p \mapsto \psi_{r+1,p}]$: $\psi_{r+1,q}$ is the same as $\psi_{r,q}$ except that $\psi_{r+1,p}$ replaces x_p. Hence, the two partial games are identical until a move to some (n, x_p) in the first game and a corresponding move to $(\mathsf{n}, \psi^k_{r+1,p})$ (for some $k \leq l$) in the second one. In such case the first game stops whereas the second one moves deterministically to $(\mathsf{n}, \psi^k_{r,p})$ and stops there as well (the index k depends on which occurrence of x_p was reached in the game on the left, but we do not care about it).

We get the second premise (2.17) from:

$$\begin{aligned}\mathcal{G}(\mathcal{A}), (\mathsf{m}, p) &\bowtie_{S} \mathcal{G}(\psi_{r+1,p}), (\mathsf{m}, \psi_{r,p}) \\ &\bowtie_{\mathsf{Id}(\mathcal{S}')} \mathcal{G}(\psi_{r+1,q}), (\mathsf{m}, \psi^k_{r,p})\end{aligned}$$

because $S \circ \mathsf{Id}(\mathcal{S}') = S$. The first line is (2.14). For the second one observe that once the game moves to $\psi^k_{r,p}$ for some particular k it cannot leave $\mathsf{SubFor}(\psi^k_{r+1,p})$. A move leaving $\mathsf{SubFor}(\psi^k_{r+1,p})$ would only be possible if there was a variable free in $\psi_{r+1,p}$ but bound in its proper superformula (and hence also bound in $\psi_{r+1,q}$). However, if x_s is bound in $\psi_{r+1,q}$ then $x_s \notin \mathsf{Var}_{r+1\leq}$, whereas it can be free in $\psi_{r+1,p}$ only if $x_s \in \mathsf{Var}_{r+1\leq}$. Thus, the reachable parts of the two compared games are isomorphic. $\square$

Chapter 3

Model Theory for ML

In this chapter we ask classical model-theoretic questions in the modal context. While model theory for single modal formulae is arguably well-understood, the theory for entire *sets* of formulae turns out to be surprisingly underdeveloped. The work presented here aims at filling this gap. It splits into two topics: *bisimulational categoricity* and *ordinal models*, addressed in two respective sections. The question about bisimulational categoricity over ordinal models, which lies in the intersection of the two topics, is addressed in the second subsection.

Bisimulational categoricity. One of the central notions of classical model theory is that of *categoricity.* A set of sentences is called *categorical* if it has a unique model *up to isomorphism.* In the context of modal logic, bisimilarity seems more appropriate than the isomorphism. One may therefore ask about *bisimulational categoricity*: the property of having a unique model *up to bisimulation.* Note that due to the obvious limitations imposed by the Skolem-Löwenheim Theorem, the classical notion of categoricity of first-order theories is only interesting when models of fixed cardinality are considered. However, unlike with isomorphism, structures of different sizes may still be bisimilar. Thus, there is no need to relativize *bisimulational* categoricity.

It turns out that the notion of bisimulational categoricity for theories expressed in modal logic is indeed well-behaved and can be characterized in terms of image-finiteness. In Theorem 3.1.4 we show that a modal theory (i.e. a maximal consistent set of modal formulae) has a unique model up to bisimulation iff it has an image-finite model, i.e. a model where every point has finitely many $\xrightarrow{\mathsf{a}}$-successors for each $\mathsf{a} \in \mathsf{Act}$. While the right-to-left implication is (an easy folklore strengthening of) the well-known Hennessy-Milner Theorem [16], the left-to-right one requires adaptation of some classical model-theoretic tools and a simple topological argument. As such, our characterization can be thought of both as a completion of the Hennessy-Milner Theorem and as a modal version of the Ryll-Nardzewski Theorem (proven independently by Ryll-Nardzewski [27] Svenonius [31] and

Engeler [14]) which says that a maximal consistent set of FO-sentences has a unique countable model (i.e. all countable models of T are isomorphic) iff it has a model $\mathcal{M}$ where for every $d < \omega$ there could be only finitely many d-tuples of elements of $\mathcal{M}$ satisfying pairwise different FO-formulae.

A natural direction for further studies is to investigate modal logic over a fixed class $\mathcal{C}$ of models (i.e. replace "model" with "model from $\mathcal{C}$" in all the definitions). We study several such classes. First, consider *two-way* models $\mathcal{C}_{\leftrightarrows}$ (also known as bidirectional models, e.g. in [2]). These are models with two accessibility relations $\to$ and $\to^{-1}$ such that one is the converse of the other. Theorem 3.1.5 says that over $\mathcal{C}_{\leftrightarrows}$ categoricity is equivalent to having a model with finite in- and out-degree. Another interesting case is the class $\mathcal{C}_{\to^+}$ of all *transitive* models: monomodal models where the only accessibility relation is transitive. Here the situation is a bit more subtle. Theorem 3.1.6 states that under an additional assumption of finiteness of Prop an analogous characterization is true: bisimulational categoricity is equivalent to having a finite model.

Furthermore, Example 3.1.20 demonstrates that with $\mathcal{C}_{\to^+}$ the additional assumption is necessary. Example 3.1.21 goes even further: over the class $\mathcal{C}_{\forall}$ of all *universal models*, i.e. ones with one arbitrary accessibility relation and another one linking every two points, an analogous characterization is false, even with no atomic propositions whatsoever. This also shows that FO-definability of the class $\mathcal{C}$ under consideration, although often helpful, is not always sufficient for a characterization similar to the ones we have. The last characterization, Theorem 3.2.3 concerning *ordinal models* $\mathcal{C}_{\mathsf{Ord}}$, demonstrates that simplicity of the considered class is not necessary either. Since the class $\mathcal{C}_{\mathsf{Ord}}$ behaves differently from all the previous ones, we investigate it separately in the second section.

Ordinal models. The second section of this chapter is devoted to the study of ordinal models $\mathcal{C}_{\mathsf{Ord}}$: monomodal models whose only accessibility relation is a (descending and strict) linear well-founded order $>$ on the universe. We start with Theorem 3.2.3 which says that within such ordinal models bisimulational categoricity is the same as having a finite model. The characterization is similar to the previous ones but the proof, although not hard, is different.

Further in the section we address another fundamental property: compactness. If every finite fragment of a set t of formulae is satisfiable in $\mathcal{C}$, does it have to be the case for the entire t? For FO-definable (or even FO-axiomatizable) classes of models, such as all the mentioned classes other than $\mathcal{C}_{\mathsf{Ord}}$, compactness of ML is a straightforward consequence of compactness of FO. $\mathcal{C}_{\mathsf{Ord}}$, however, cannot be axiomatized in FO and in fact FO is not compact over $\mathcal{C}_{\mathsf{Ord}}$. This makes the question about compactness of ML over $\mathcal{C}_{\mathsf{Ord}}$ nontrivial. The Example 3.2.5 showing the lack of compactness with infinite Prop is rather straightforward. The complementary Theorem 3.2.6, however, has a more involved proof. It says that, perhaps surprisingly, if Prop is finite then ML *is* compact over $\mathcal{C}_{\mathsf{Ord}}$.

We conclude the chapter with Theorem 3.2.12 which expresses a short model property: every set of modal formulae satisfiable in $\mathcal{C}_{\mathsf{Ord}}$ has an ordinal model that is not too long. The ordinal bound on the length of such a model depends on the number of colors and is shown to be strict in Example 3.2.11.

History and credits. The contribution of the first section, although phrased in slightly different terms, was published in [22] which continued the line of research initiated in the author's master thesis [21]. The only significant exception is Example 3.1.16 which is a counterexample to what was initially conjectured. The results from the second section are, to the best of our knowledge, all new.

3.1 Bisimulational Categoricity

3.1.1 Introduction

Let us start with recalling the tight connection between modal logic, first-order logic and bisimulation, given by the famous van Benthem Theorem. Although we will mostly only need the easier of its two implications, for the sake of completeness we present the full theorem as an important element of the wider picture. The characterization identifies ML as precisely the fragment of FO that is invariant under bisimulation, exposing a deep link between modal logic and bisimilarity.

Consider the *standard translation*, which is an embedding $\mathsf{ST}_x(\_) : \mathsf{ML} \to \mathsf{FO}$ that maps every modal sentence φ to an FO formula $\mathsf{ST}_x(\varphi)$ with one free variable x:

$$\begin{aligned}
\mathsf{ST}_x(\top) &= \top \quad \text{and} \quad \mathsf{ST}_x(\bot) = \bot \\
\mathsf{ST}_x(\tau) &= \tau(x) \quad \text{and} \quad \mathsf{ST}_x(\neg\tau) = \neg\tau(x) \\
\mathsf{ST}_x(\varphi \vee \psi) &= \mathsf{ST}_x(\varphi) \vee \mathsf{ST}_x(\psi) \quad \text{and} \quad \mathsf{ST}_x(\varphi \wedge \psi) = \mathsf{ST}_x(\varphi) \wedge \mathsf{ST}_x(\psi) \\
\mathsf{ST}_x(\langle \mathsf{a} \rangle \varphi) &= \exists_{x \xrightarrow{\mathsf{a}} y}.\mathsf{ST}_y(\varphi) \quad \text{and} \quad \mathsf{ST}_x([\mathsf{a}]\varphi) = \forall_{x \xrightarrow{\mathsf{a}} y}.\mathsf{ST}_y(\varphi)
\end{aligned}$$

with $x \neq y$. The FO formula uses the same modal signature as the original φ, i.e. the signature with a binary symbol $\xrightarrow{\mathsf{a}}$ for each $\mathsf{a} \in \mathsf{Act}$ and a unary τ for each $\tau \in \mathsf{Prop}$. The standard translation of $\langle \mathsf{a} \rangle \varphi$ reflects the semantics of $\langle \mathsf{a} \rangle$ as a restricted form of existential quantification: $\langle \mathsf{a} \rangle \varphi$ means "there exists an $\xrightarrow{\mathsf{a}}$-child of the current point that satisfies φ". Symmetrically, $[\mathsf{a}]$ is a restricted universal quantifier. The above translation is designed so that:

$$\mathcal{M}, \mathsf{m} \models \varphi \iff \mathcal{M} \text{ satisfies } \mathsf{ST}_x(\varphi)(\mathsf{m})$$

for every model $\mathcal{M}$ and point $\mathsf{m} \in M$. Of course, it follows from invariance of ML under bisimulation that the image of $\mathsf{ST}_x(\_)$ is invariant under bisimulation as well. The famous van Benthem Theorem [33, Theorem 1.9] says that also the converse holds: ML is *precisely* the bisimulation invariant fragment of FO.

Theorem 3.1.1 (van Benthem)**.** *For every $\varphi(x) \in \mathsf{FO}$, φ is invariant under bisimulation $\iff$ it is equivalent to the standard translation of some modal formula.*

Here, invariance under bisimulation of an FO formula $\varphi(x)$ with one free variable x means that $\mathcal{M}, \mathsf{m} \leftrightarrows \mathcal{M}', \mathsf{m}'$ implies that $\varphi(\mathsf{m})$ in $\mathcal{M}$ iff $\varphi(\mathsf{m}')$ in $\mathcal{M}'$.

Given how natural the notion of a bisimulation is and how closely it is related to modal logic, it is also natural to ask when modal logic can describe a (pointed) model *uniquely*. That is, given a set of modal sentences, when does it happen that it has a unique model *up to bisimulation*? A reader familiar with model theory of first-order logic may realize that this question is analogous to the notion of *categoricity*, the property of having a unique model *up to isomorphism*.

Invariance of modal logic under bisimulation means that if points m and m' are *bisimilar* then they are always *modally equivalent* (i.e. satisfy exactly the same formulae, which we denote by $\mathsf{m} \equiv_{\mathsf{ML}} \mathsf{m}'$). Therefore, the quest for characterizing bisimulational categoricity boils down to the question about conditions under which the converse implication holds. As mentioned, in *image-finite* models points that are logically indistinguishable have to be bisimilar. The following example shows that without the assumption of image-finiteness this does not have to be the case.

Example 3.1.2. The Hedgehogs: $\mathcal{H}, \mathsf{root}_{\mathcal{H}}$ and $\mathcal{H}', \mathsf{root}_{\mathcal{H}'}$:

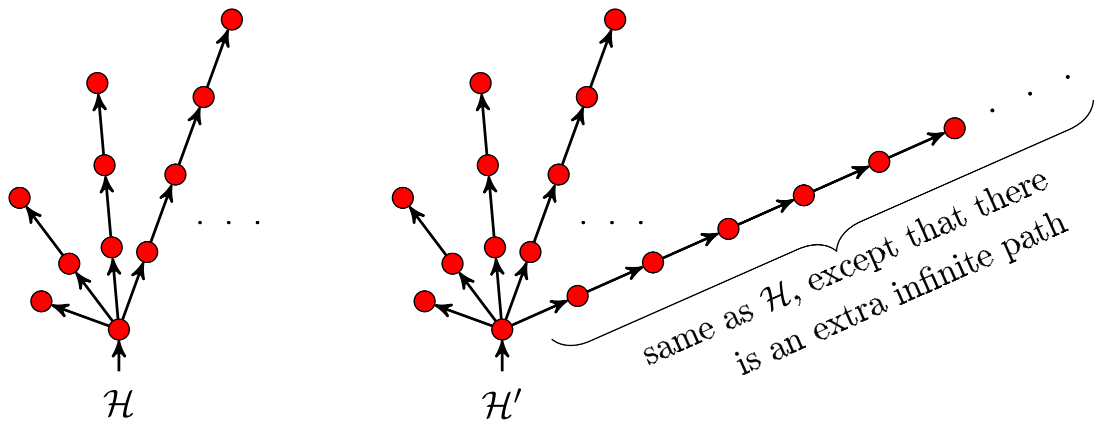

The two models are not bisimilar, as one of them is well-founded but the other is not. However, it is easy to show that they cannot be distinguished by ML formulae. In fact, even the full first-order logic cannot distinguish the models, as can be shown using Ehrenfeucht-Fraïssé games (these are a classical tool characterizing expressive power of FO similarly to depth-k bisimilarity $\leftrightarrows^k$ characterizing expressive power of ML, see [24] for a reference).

As it turns out, the above example is an illustration of a general phenomenon, which is that among infinitely many behaviors one can always find a *limit* one that: (i) can be either included or removed from the model but (ii) our local logical means are too weak to tell the difference. This will be the key intuition underlying our characterization of bisimulational categoricity, which says that the requirement of image-finiteness is not only sufficient, but also necessary.

In order to formulate the theorem, we first formally introduce the notion of a *type*, i.e. a maximal consistent set of formulae, analogous to types in first-order model theory (here by *type* we always mean a *complete* one).

Definition 3.1.3. Given a point m in a model $\mathcal{M}$, its *modal type* or *modal theory*, denoted $\mathbf{tp}^{\mathcal{M}}(\mathsf{m})$, is the set $\{\varphi \in \mathsf{ML} \mid \mathsf{m} \in [\![\varphi]\!]^{\mathcal{M}}\}$ of all modal formulae that it satisfies. The set of all modal types will be denoted $\mathbb{T}$. Sometimes we will be only interested in models from a fixed class $\mathcal{C}$. In such case $\mathbb{T}_{\mathcal{C}}$ will denote the set of all the types that are present *in some model from* $\mathcal{C}$.

We can now formulate our first characterization. By definition, every type has a model. The following theorem tells when such a model is *unique*.

Theorem 3.1.4. *For every type $t \in \mathbb{T}$, the following are equivalent:*

(1) *t has a unique model up to $\leftrightarrows$;*

(2) *every model of t is bisimilar to an image-finite model;*

(3) *t has a model which is image-finite.*

Note that although the above theorem describes bisimulational categoricity over the class of all models, this does not automatically yield an analogous characterization for its arbitrary subclasses. For instance, if we consider the class $\mathcal{C}_{\mathsf{WF}}$ of all *well-founded* models then the Hedgehog $\mathcal{H}$ of Example 3.1.2 is a unique (up to bisimulation) model of its type, despite being inherently image-infinite (the reasons for this are discussed in Subsection 3.1.3). Nevertheless, we provide characterizations analogous to Theorem 3.1.4 for some interesting classes of models.

A *two-way* model (also called *bidirectional* model e.g. in [2]) is a model over signature consisting of two accessibility relations, denoted $\to$ and $\to^{-1}$, such that one is the reverse of the other (meaning that $\mathsf{m} \to \mathsf{n}$ iff $\mathsf{n} \to^{-1} \mathsf{m}$). We denote the class of all two-way models by $\mathcal{C}_{\leftrightarrows}$. Then:

Theorem 3.1.5. *For every type $t \in \mathbb{T}_{\mathcal{C}_{\leftrightarrows}}$, the following are equivalent:*

(1) *t has a two-way model which is, up to $\leftrightarrows$, unique among all two-way models;*

(2) *every two-way model of t is bisimilar to an image-finite two-way model;*

(3) *t has an image-finite two-way model.*

Consider the class $\mathcal{C}_{\to^+}$ of all *transitive* models, i.e. models with a single accessibility relation $\to^+$ that is required to be transitive (in the context of $\mathcal{C}_{\to^+}$ we will use the equivalent expressions *descendant* and *child* interchangeably; similarly with *ancestor* and *parent*). Then:

Theorem 3.1.6. *Assume that the set* Prop *of atomic propositions is finite. For every type $t \in \mathbb{T}_{\to^+}$, the following are equivalent:*

(1) *t has a transitive model which is, up to $\leftrightarrows$, unique among all transitive models;*

(2) *every transitive model of t is bisimilar to a finite transitive model;*

(3) *t has a finite transitive model.*

Note that for transitive models image-finiteness is the same as finiteness, assuming that all the points are accessible from the root. Thus, in light of Proposition 3.1.17 which says that in $\mathcal{C}_{\to^+}$ every finite model can be uniquely described with a single modal sentence, the last theorem implies that when it comes to defining transitive models up to bisimulation, the expressive power of modal logic does not increase when we move from single formulae to entire theories.

In contrast to the above theorems, in Subsection 3.1.3 we will show counterexamples limiting possible extensions of our characterizations:

- Example 3.1.20 shows that, perhaps surprisingly, the assumption of finiteness of Prop in Theorem 3.1.6 is necessary.
- Example 3.1.21 demonstrates that even with finite vocabulary, compactness of the logic over the class $\mathcal{C}$ (or even FO-axiomatizability of $\mathcal{C}$, which is stronger than compactness) does not guarantee an analogous characterization. We consider the class $\mathcal{C}_\forall$ of all *universal* models, i.e. models with two relations $\to$ and $\to^\forall$ with the second one being always full (i.e. linking every two points). We then show a complete theory whose model is unique but infinite (which means that every point has infinitely many $\to^\forall$-successors).

Classes of models vs altered semantics. Considering various classes of models and investigating model-theoretic questions related to ML over these classes is a neat, uniform way of defining the subject of our study. However, an important motivation for such framework comes from a different point of view, where a class $\mathcal{C}$ is thought of as encoding modal logic interpreted over *all* models but with modal operators having *altered* semantics.

- Consider *two-way* modal logic, i.e. monomodal logic enriched with backward modalities: $\Diamond^{-1}\varphi$ interpreted as "there exists a *predecessor* satisfying φ". Analogously, we obtain the notion of a two-way bisimulation by extending the standard Definition 2.1.4 with additional *backward* back and forth conditions in which the term "predecessor" replaces "successor". Then, the notion of a two-way model allows us to capture such two-way semantics without changing the definitions of modal logic or bisimulation. This is because going *backward* along an edge $\to$ is the same as going *forward* along the opposite edge $\to^{-1}$. Given a monomodal model $\mathcal{M}$, denote by $\mathcal{M}_{\leftrightarrows}$ the bimodal extension of $\mathcal{M}$ with the same universe and interpretation of $\to$ and $\to^{-1}$ interpreted as the converse of $\to$. It follows that $\mathsf{m}, \mathsf{n} \in M$ are (i) equivalent with respect to the two-way ML and (ii) two-way bisimilar in $\mathcal{M}$ iff they are (i) modally equivalent and (ii) bisimilar in $\mathcal{M}_{\leftrightarrows}$, respectively. Moreover, $\mathcal{M}_{\leftrightarrows}$ is image-finite iff all the points in $\mathcal{M}$ have finite in- and out-degree. Thus, Theorem 3.1.5 can be reformulated as: a complete theory in two-way ML has a unique model up to two-way bisimulation iff it has a model where every point has finite in- and out-degree.

- Similarly, one can consider *transitive* modal logic and *transitive* bisimulations (also known as the *EF-logic* and *EF-bisimulations*, respectively [4]) where "descendant" replaces "successor" in all the definitions. Then, Theorem 3.1.6 says that a complete theory in the transitive modal logic has a model unique up to transitive bisimulation iff it has a finite model. To see that the alternative formulation is equivalent, for each monomodal $\mathcal{M}$ take $\mathcal{M}_{\to^+}$ whose accessibility relation is the transitive closure of that from $\mathcal{M}$. Then, equivalence with respect to transitive ML and transitive bisimilarity in $\mathcal{M}$ are the same as ordinary modal equivalence and bisimilarity in $\mathcal{M}_{\to^+}$, respectively.

- In the case of $\mathcal{C}_\forall$, one can equivalently consider modal logic enriched with *universal* modalities ($\langle\exists\rangle\varphi$ interpreted as "there exists a point satisfying φ") and global bisimulations (which are bisimulations $Z \subseteq M \times M'$ whose projections on both coordinates are equal to the full universes $\pi_1(Z) = M$ and $\pi_2(Z) = M'$). Then our Example 3.1.21 is a model of such universal modal logic that is infinite yet unique up to global bisimulation.

3.1.2 Proofs of the Characterizations

In all the three cases, the implication (2) $\Rightarrow$ (3) is immediate, as by definition every type has a model and modal logic is invariant under bisimulation.

The implication (3) $\Rightarrow$ (1) follows from a generalization of the Hennessy-Milner Theorem [16]:

Proposition 3.1.7 (à la Hennessy-Milner)**.** *Assume an image-finite model $\mathcal{M}$. Then, for every $\mathcal{M}'$ and every $\mathsf{m} \in \mathcal{M}, \mathsf{m}' \in \mathcal{M}'$:*

$$\mathcal{M}, \mathsf{m} \equiv_{\mathsf{ML}} \mathcal{M}', \mathsf{m}' \qquad \underline{\textit{implies}} \qquad \mathcal{M}, \mathsf{m} \leftrightarroows \mathcal{M}', \mathsf{m}'.$$

This is a well-known folklore strengthening of the classical Hennessy-Milner Theorem, where both $\mathcal{M}$ and $\mathcal{M}'$ are assumed to be image-finite. Instead of proving it with bare hands, which is not difficult, let us introduce a higher-level point of view that will be important later. Along the same lines as in the classical model theory for first-order logic (as found e.g. in [24]), our types can be equipped with a topology turning it into a Hausdorff space.

Definition 3.1.8. For every $\varphi \in \mathsf{ML}$, we take the set $\langle\varphi\rangle = \{t \in \mathbb{T}_{\mathcal{C}} \mid \varphi \in t\}$ of all types containing it. Then, the set $\{\langle\varphi\rangle \mid \varphi \in \mathsf{ML}\}$ is a basis of clopen sets generating a topology on $\mathbb{T}$.

Alternatively, in the case when the signature is at most countable, one could obtain the same topology by first fixing any enumeration of ML formulae and then defining a metric $d(t, t') = \frac{1}{n}$ for n being the number of the first formula on which t and t' differ (and 0 if $t = t'$). The underlying intuition is that types which are similar, i.e. hard to distinguish, should be close to each other.

Proposition 3.1.9. *Analogously to the first-order case [24], we have that for every class of models $\mathcal{C}$:*

- *the space $\mathbb{T}_{\mathcal{C}}$ is Hausdorff;*
- *the logic ML is compact over $\mathcal{C}$ (i.e. if every finite fragment of a set of formulae t is satisfiable in $\mathcal{C}$, then so is the entire t) $\iff$ the space $\mathbb{T}_{\mathcal{C}}$ is compact;*
- *given $T \subseteq \mathbb{T}_{\mathcal{C}}$, $t \in \mathbb{T}_{\mathcal{C}}$ is isolated in T $\iff$ there exists a single ML formula $\varphi \in t$ s.t. $\varphi \notin t'$ for every other $t' \in T$.*

Proof. Observe that by identifying a type with its characteristic function we can view the space $\mathbb{T}_{\mathcal{C}}$ as a subspace of 2^{ML}. Since the later is Hausdorff, so is $\mathbb{T}_{\mathcal{C}}$. Moreover, a subspace of a compact Hausdorff space is compact iff it is closed, and it is easy to check that closedness of $\mathbb{T}$ is the same as logical compactness of ML over $\mathcal{C}$. The last item follows from the observation that in any topological space, a point is isolated iff it is isolated by a basic open set. $\square$

Let us now prove Proposition 3.1.7:

Proof. It suffices to show that the relation $\equiv_{\mathsf{ML}} \subseteq M \times M'$ of modal equivalence is itself a bisimulation. The base condition is immediate.

For the back and the forth conditions, let us take $\mathsf{n} \equiv_{\mathsf{ML}} \mathsf{n}'$, and any $\mathsf{a} \in \mathsf{Act}$. By our assumption, n can only have a finite number of $\xrightarrow{\mathsf{a}}$-children. In particular, they have only a finite number of distinct modal types $t_1, ..., t_k \in \mathbb{T}$. Since $\mathbb{T}$ is a Hausdorff space, we can find pairwise disjoint basic open neighborhoods $\langle\varphi_1\rangle, ..., \langle\varphi_k\rangle$ of these types, which by definition means that there are mutually exclusive formulae $\varphi_1, ..., \varphi_k$ with $\varphi_i \in t_i$ but $\varphi_i \notin t_j$ for all $i \neq j$. Both n, and by equivalence also n', satisfy:

$$[\mathsf{a}](\bigvee_{i \in \{1,...,k\}} \varphi_i); \qquad \bigwedge_{i \in \{1,...,k\}} \langle \mathsf{a} \rangle \varphi_i; \qquad \{[\mathsf{a}](\varphi_i \Rightarrow \psi) \mid \psi \in t_i\}.$$

It follows that the types of $\xrightarrow{\mathsf{a}}$-children of n' are exactly $t_1, ..., t_k$. But this implies both the forth and the back conditions, as it means that for every $\xrightarrow{\mathsf{a}}$-child of n (or n', respectively) there exists an equivalent $\xrightarrow{\mathsf{a}}$-child of n' (n, respectively). □

Towards the last (and hardest to prove) implication (1) $\Rightarrow$ (2), let us establish a few facts about modal logic.

An important notion is that of modal saturation (also called m-saturation). Our topology on types allows us to capture it in an elegant way.

Definition 3.1.10. We say that a point m in a model $\mathcal{M}$ is *modally saturated* if for every $\mathsf{a} \in \mathsf{Act}$, the set of types of its $\xrightarrow{\mathsf{a}}$-children $\{\mathbf{tp}^{\mathcal{M}}(\mathsf{n}) \mid \mathsf{m} \xrightarrow{\mathsf{a}} \mathsf{n}\}$ is a closed subset of $\mathbb{T}$. We call $\mathcal{M}$ *modally saturated* if all its points are modally saturated.

In more concrete terms (the way modal saturation is usually defined [2]): if every finite fragment of t is realized in some $\xrightarrow{\mathsf{a}}$-child of m, then there exists an m's $\xrightarrow{\mathsf{a}}$-child realizing the entire t.

Note that it is immediate that modal saturation generalizes the notion of image-finiteness, as in a Hausdorff space finite sets are always closed.

The following classical fact says that for saturated models modal equivalence is the same as bisimilarity.

Proposition 3.1.11. *Given any two modally saturated models $\mathcal{M}, \mathcal{M}'$:*

$$\mathcal{M}, \mathsf{m} \equiv_{\mathsf{ML}} \mathcal{M}', \mathsf{m}' \qquad \underline{\textit{implies}} \qquad \mathcal{M}, \mathsf{m} \leftrightarrows \mathcal{M}', \mathsf{m}'$$

for any $\mathsf{m} \in \mathcal{M}, \mathsf{m}' \in \mathcal{M}'$.

Remark 3.1.12. A natural question in our context is whether the van Benthem theorem is true if we only take models from a fixed class $\mathcal{C}$ into account. The standard translation $\mathsf{ST}_x(\_)$ produces a first-order formula that is equivalent to the original one in every model and invariant under bisimulation.

However, the more interesting direction is the opposite one, i.e. given an FO formula invariant under bisimulation, to find an equivalent ML formula. Here, an FO formula $\varphi(x)$ is invariant under bisimulation *over the class* $\mathcal{C}$ if for every $\mathcal{M}$ and $\mathcal{M}'$ *both belonging to* $\mathcal{C}$: $\mathcal{M}, \mathsf{m} \leftrightarrows \mathcal{M}', \mathsf{m}'$ implies that $\varphi(\mathsf{m})$ in $\mathcal{M}$ iff $\varphi(\mathsf{m}')$ in $\mathcal{M}'$. Note that restricting the class $\mathcal{C}$ weakens the condition, which possibly makes *more* formulae invariant under bisimulation. Finding a modal equivalent of a formula φ that is bisimulation-invariant over $\mathcal{C}$ is a nontrivial task and sometimes it is just not possible, as illustrated by the class:

$$\mathcal{C} = \mathcal{C}_{\to^+} \cap \mathcal{C}_{\text{fin}}$$

of models that are both transitive and finite. Finiteness of a model $\mathcal{M}$ implies that a $\mathsf{m} \in \mathcal{M}$ is not well-founded (an infinite path starts there) iff m belongs to a cycle $\mathsf{m} \to \dots \to \mathsf{m}$. In turn, transitivity implies that this property is equivalent to the existence of a self-loop $\mathsf{m} \to \mathsf{m}$. The property can be expressed with the FO-formula $\varphi(x) = x \to x$ and since well-foundedness is invariant under bisimulation, so is φ over $\mathcal{C}$. However, no modal formula can be equivalent to φ. To see this, consider a pair of models $\mathcal{M}_k$ and $\mathcal{M}'_k$ from $\mathcal{C}$ for every $k < \omega$:

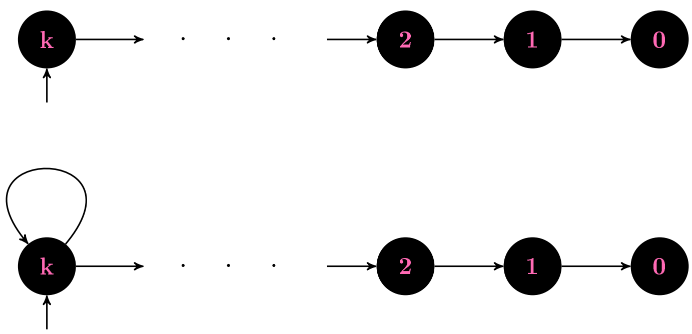

As usual with transitive models, for clarity we skip some arrows belonging to the transitive closure of the depicted ones. No modal formula of modal-propositional depth at most k can differentiate between $\mathcal{M}_k$ and $\mathcal{M}'_k$. Since only $\mathcal{M}_k$ but not $\mathcal{M}'_k$ is well-founded, this implies that well-foundedness over $\mathcal{C}$ cannot be captured with ML. The equivalence $\mathcal{M}_k, k \equiv^k_{\mathsf{ML}} \mathcal{M}'_k, k$ follows from the game characterization (2.8). In order to win the k-round game, it suffices if $\exists$ve preserves the following invariant: either positions in both models are equal, or both are greater than the remaining number of rounds.

Still, the van Benthem characterization is true in many interesting classes of models, such as all finite models, or all well-founded ones, but the proof often requires new ideas (see [13] for a good summary). Here, let us only state a much easier result, which essentially follows from the proof of the classical version of the theorem (as found in [2]).

Assume that a class $\mathcal{C}$ of models is axiomatized by a set A of FO-sentences (meaning that $\mathcal{C}$ is the class of models satisfying A). Then, ML is precisely the fragment of FO that is invariant under bisimulation *over* $\mathcal{C}$:

Theorem 3.1.13 (Relativized van Benthem)**.** *Assume that $\mathcal{C}$ is axiomatizable in FO. Then, for every $\varphi(x) \in \mathsf{FO}$: φ is invariant under bisimulation over $\mathcal{C}$ $\iff$ it is equivalent, in all models from $\mathcal{C}$, to the standard translation of a modal formula.*

All the classes $\mathcal{C}_{\leftrightharpoons}$, $\mathcal{C}_{+}$ and $\mathcal{C}_{\forall}$ (and of course the class of all models) can be described by FO sentences: $\forall_{x,y}.x \to y \iff y \to^{-1} x$; $\forall_{x,y,z}.x \to^{+} y \wedge y \to^{+} z \implies x \to^{+} z$ and $\forall_{x,y}.x \to^{\forall} y$, respectively. Thus, the theorem applies.

The following fact uses (the easier implication of, i.e. the correctness of the standard translation) the above theorem:

Proposition 3.1.14. *Assume an FO-axiomatizable class of models $\mathcal{C}$. Then:*

1. *ML over $\mathcal{C}$ is compact, meaning that if every finite fragment of a set $t \subseteq \mathsf{ML}$ has a model in $\mathcal{C}$ then so does the entire t;*

2. *ML over $\mathcal{C}$ has the modally saturated model property, meaning that every $t \subseteq \mathsf{ML}$ satisfiable in $\mathcal{C}$ has a modally saturated model in $\mathcal{C}$*

Proof. Let A be the set of FO formulae axiomatizing $\mathcal{C}$.

For the first item, observe that t is satisfiable in $\mathcal{C}$ iff $A \cup \mathsf{ST}_x[t]$ is satisfiable (in any model). Hence, compactness of ML over $\mathcal{C}$ follows from compactness of FO.

For the second item, take an ω-saturated model of $A \cup \mathsf{ST}_x[t]$. Since this set of formulae is satisfiable, such a model must exist, and it is straightforward to check that ω-saturated models are also modally saturated ([24] is a good reference for the basics of FO model theory). $\square$

Let us recall an elementary topological fact which, although simple, is the heart of our characterizations. Since any infinite compact space contains a non-isolated point and closed subspaces of a compact space are always compact, it follows that:

Lemma 3.1.15. *If Y is a closed infinite subset of a compact topological space X, then it contains a point $y \in Y$ that is not isolated in Y.*

We are ready to prove the only missing implication (1) $\Rightarrow$ (2) in all the three theorems: Theorem 3.1.4 about all models, Theorem 3.1.5 about two-way models and Theorem 3.1.6 about the transitive ones. Let us start with the class of all models.

All Models

Take a model $\mathcal{M}, \mathsf{m}$ that is not bisimilar to any image-finite model. We will construct another model that is equivalent but non-bisimilar to it. We may combine: (i) Proposition 3.1.14 to obtain an equivalent model which is modally saturated, (ii) Proposition 2.1.8 to take its quotient by $\leftrightarrows$ where (by Proposition 3.1.11) no two points satisfy the same formulae and finally (iii) Proposition 2.1.6 to take a submodel accessible from the root. If such model is not bisimilar to $\mathcal{M}, \mathsf{m}$ then we are done, so the remaining case is when $\mathcal{M}, \mathsf{m}$ has all the properties listed above.

Since by our assumption $\mathcal{M}, \mathsf{m}$ is not image-finite, there must exist a point n reachable from m by a finite path and having infinitely many $\xrightarrow{\mathsf{a}}$-children for some $\mathsf{a} \in \mathsf{Act}$. The set $T = \{\mathbf{tp}(\mathsf{n}') \mid \mathsf{n} \xrightarrow{\mathsf{a}} \mathsf{n}'\}$ is an infinite subset of the compact space $\mathbb{T}$. Modal saturation of $\mathcal{M}$ means that T is closed, and so by Lemma 3.1.15 it contains a non-isolated limit type $t^{\lim}$ realized in some $\xrightarrow{\mathsf{a}}$-child $\mathsf{n}^{\lim}$ of n. Now, in order to construct another model for t we take $\mathcal{N}$ that is identical to $\mathcal{M}$ except that:

$$\xrightarrow{\mathsf{a}}^{\mathcal{N}} = \xrightarrow{\mathsf{a}}^{\mathcal{M}} - \{(\mathsf{n}, \mathsf{n}^{\lim})\},$$

i.e. we remove the arrow $\xrightarrow{\mathsf{a}}$ leading from n to $\mathsf{n}^{\lim}$.

We prove by induction on $k < \omega$ that any point $\mathsf{n} \in \mathcal{M}$ satisfies exactly the same formulae of modal depth k in both $\mathcal{M}$ and $\mathcal{N}$ (and thus in particular $\mathcal{N}, \mathsf{m} \models t$). The base case is obvious. For the induction step, the only interesting case is for n, as prima facie it could satisfy fewer sentences of the form $\Diamond\varphi$. However, since $t^{\lim}$ is not isolated in T, for any $\varphi \in t^{\lim}$ there must be $t' \in T$ s.t. $\varphi \in t'$. By definition of T this means that there is a sibling n' of $\mathsf{n}^{\lim}$ such that $\mathcal{M}, \mathsf{n}' \models t'$. In particular, $\mathcal{M}, \mathsf{n}' \models \varphi$. But modal depth of φ is smaller than that of $\Diamond\varphi$, so we know by induction hypothesis that $\mathcal{N}, \mathsf{n}' \models \varphi$ and hence $\mathcal{N}, \mathsf{n} \models \Diamond\varphi$.

On the other hand, we will show that $\mathcal{M}, \mathsf{m} \not\leftrightarrows \mathcal{N}, \mathsf{m}$, as $\forall$dam has the following winning strategy in the bisimulation game: (i) First follow the path to the point n in $\mathcal{M}$. If after that $\exists$ve responds with a point $\mathsf{n}' \in \mathcal{N}$ other than n, we know that $\mathcal{M}, \mathsf{n} \not\equiv_{\mathsf{ML}} \mathcal{N}, \mathsf{n}'$ (as no two different points are equivalent in $\mathcal{N}$) and so $\mathcal{M}, \mathsf{n} \not\leftrightarrows \mathcal{N}, \mathsf{n}'$, which means that $\forall$dam can now win the game. (ii) If $\exists$ve responded with the same point $\mathsf{n} \in \mathcal{N}$, $\forall$dam moves to $\mathsf{n}^{\lim}$ in $\mathcal{M}$. Now $\exists$ve has to respond with some point $\mathsf{n}' \in \mathcal{N}$ but by definition of $\mathcal{N}$ she cannot choose $\mathsf{n}^{\lim}$ and so again $\mathcal{M}, \mathsf{n}^{\lim} \not\equiv_{\mathsf{ML}} \mathcal{N}, \mathsf{n}'$, meaning that $\forall$dam can win the game from that point. This completes the proof of Theorem 3.1.4.

Two-Way Models

In the case of all two-way models, we need a slight modification of the previous construction due to the fact that in such models to remove an arrow

$\mathsf{m} \to \mathsf{n}$ we also have to remove the opposite one $\mathsf{n} \to^{-1} \mathsf{m}$.

As in the previous case, we take a modally saturated model of $t \in \mathbb{T}_{\mathcal{C}_{\leftrightarrows}}$ where any two different points have different types and any point is accessible by a finite path (possibly using both $\to$ and $\to^{-1}$) from the root. This is possible, because the class $\mathcal{C}_{\leftrightarrows}$ of two-way models is closed under both generated submodels and quotients by bisimulations. Assume that some point $\mathsf{n} \in \mathcal{M}$ has infinitely many $\to$-children (the case with infinitely many $\to^{-1}$-children is entirely symmetric) and take the limit $t^{\lim}$ of $T = \{\mathbf{tp}(\mathsf{n}') \mid \mathsf{n} \to \mathsf{n}'\}$ realized by some $\mathsf{n}^{\lim}$.

We define an equivalent but not bisimilar model as follows. First take the disjoint union $\mathcal{N}' = \mathcal{M}_1 + \mathcal{M}_2 + \mathcal{M}_3$, where each $\mathcal{M}_i$ is a copy of $\mathcal{M}$. We will denote the element of $\mathcal{M}_i$ corresponding to $\mathsf{n}' \in M$ by n'_i. Let us also pick any child $\mathsf{n}' \in M$ of n different than $\mathsf{n}^{\lim}$. Then, our model $\mathcal{N}$ is just $\mathcal{N}'$ without the arrow $\mathsf{n}_2 \to \mathsf{n}_2^{\lim}$ and with two additional arrows $\mathsf{n}_2 \to \mathsf{n}'_1$ and $\mathsf{n}_3 \to \mathsf{n}_2^{\lim}$:

$$\to^{\mathcal{N}} = \to^{\mathcal{N}'} - \{(\mathsf{n}_2, \mathsf{n}_2^{\lim})\} \cup \{(\mathsf{n}_2, \mathsf{n}'_1), (\mathsf{n}_3, \mathsf{n}_2^{\lim})\}$$

and $\to^{-1}$ modified accordingly.

A picture of $\mathcal{M}, \mathsf{m}$ and $\mathcal{N}, \mathsf{m}_1$:

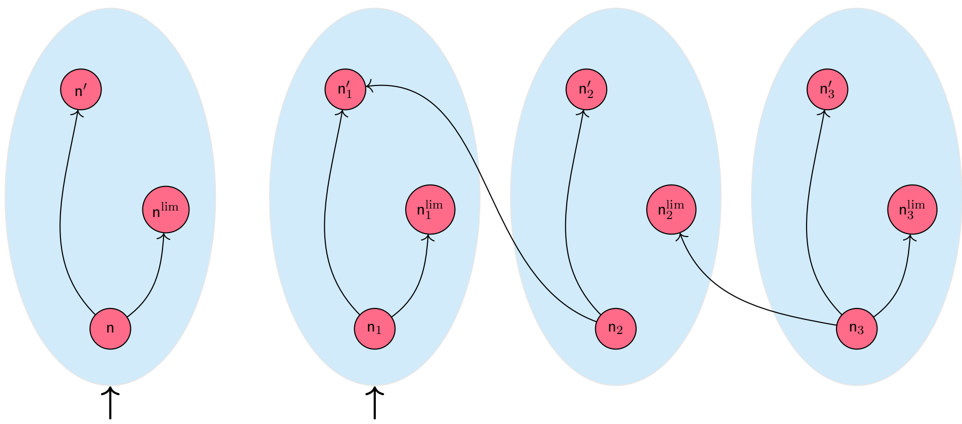

The rest of the proof is analogous to the previous case. We first prove by induction on $k < \omega$ that for every $\mathsf{m}' \in \mathcal{M}$, the models $\mathcal{M}, \mathsf{m}'$ and $\mathcal{N}, \mathsf{m}'_i$ satisfy the same ML-formulae of modal depth k. This boils down to checking several straightforward cases (the one in which we use the fact that $t_{\lim}$ was not isolated is that with n_2's $\to$-successors).

The winning strategy for $\forall$dam witnessing $\mathcal{M}, \mathsf{m} \not\equiv_{\mathsf{ML}} \mathcal{N}, \mathsf{m}_1$ is as follows: (i) First follow the path from m_1 to n_2 in $\mathcal{N}$. In order not to loose, $\exists$ve has to respond in $\mathcal{M}$ with the only point that is equivalent to n_2, namely n. (ii) Then, $\forall$dam moves to $\mathsf{n}^{\lim}$ in $\mathcal{M}$ and $\exists$ve has to respond in $\mathcal{N}$ with a point non-equivalent with it, thus loosing the game.

Note that taking three copies of the original model rather than a single one is necessary in the above construction. This is because *accessibility* means *two-way accessibility*, i.e. with respect to both $\to$ and $\to^{-1}$. Thus, after removing the arrow $\mathsf{n}_2 \to \mathsf{n}_2^{\lim}$, n_2 does not have to be accessible from m_2. In fact, it could actually happen that $\mathcal{M}, \mathsf{m} \leftrightarrows \mathcal{N}, \mathsf{m}_2$. However, we know that n_1' is accessible from m_1 and from there we can move backwards to n_2.

Transitive Models

This is the most involved case. The key difficulty is that we cannot simply remove any arrow, as its existence may be forced by transitivity and presence of other arrows. Consider the following example.

Example 3.1.16. In the model below, the rightmost blue point has a copy of ω (with the reverse order as the accessibility relation) as its children. We assume that the model is transitive and for clarity do not draw the arrows implied by transitivity.

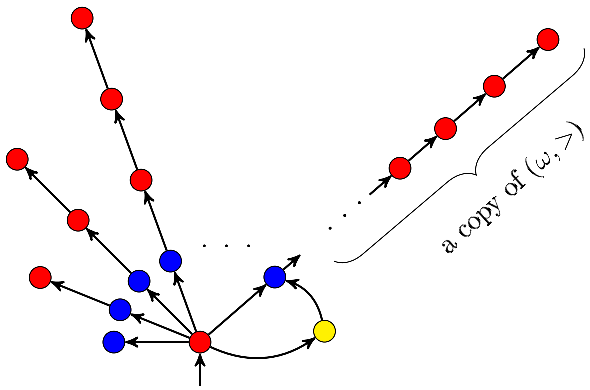

One can check that the type $t_{\lim}$ of the rightmost blue point *is not isolated* among the types of its blue siblings. However, it *is isolated* from the perspective of the yellow point, which in turn is isolated from the perspective of the root. Basing on that observation it is not hard to show that any model ML-equivalent to the one above must realize $t_{\lim}$ in a child of its root. In particular, this demonstrates that not every non-isolated type can be omitted. Nevertheless, we will show that if the set Prop of atomic propositions is finite, then in the presence of a non-isolated type it is always possible to find *some* (possibly different) type that can be omitted.

Let us start with the following simple fact:

Proposition 3.1.17. *Assume finite* Prop*. In* $\mathcal{C}_{\to^+}$*, finite models are definable up to* $\leftrightarrows$*, meaning that if* $\mathcal{M}, \mathsf{m}$ *is a finite transitive model, then there is* *a modal formula θ s.t. every transitive model of θ is bisimilar to $\mathcal{M}, \mathsf{m}$. In particular, types of finite transitive models are always isolated in $\mathbb{T}_{\mathcal{C}_{\to^+}}$.*

Proof. Since $M = \mathsf{m}_1, ..., \mathsf{m}_k$ is finite, it realizes only finitely many types $t_1, ..., t_k$ (w.l.o.g. all distinct, as otherwise we may quotient the model). Since $\mathbb{T}_{\mathcal{C}_{\to^+}}$ is a Hausdorff space, there are mutually exclusive sentences $\varphi_i \in t_i$ for every i. First, define ψ_i to be the formula that describes which atomic propositions belong to t_i and which other types it sees:

$$\bigwedge\{\tau \in \mathsf{Prop} \mid \tau \in t_i\} \wedge \bigwedge\{\neg\tau \in \mathsf{Prop} \mid \tau \notin t_i\}$$
$$\wedge$$
$$\Box(\bigvee\{\varphi_j \mid \mathsf{m}_i \to^+ \mathsf{m}_j\}) \wedge \bigwedge\{\Diamond\varphi_j \mid \mathsf{m}_i \to^+ \mathsf{m}_j\}.$$

Then, we put:

$$\theta_i = \psi_i \wedge \Box(\bigwedge_{j \in \{1,...,k\}} \{\varphi_j \Rightarrow \psi_j\})$$

It is straightforward that $\theta_i \in t_i$. On the other hand, if $\mathcal{N}, \mathsf{n} \models \theta_i$, then already $\mathcal{N}, \mathsf{n} \leftrightarrows \mathcal{M}, \mathsf{m}_i$. Indeed, w.l.o.g. we may assume that all the points of such $\mathcal{N}$ are reachable from n and then it is easy to check that: (i) the types of all the points of $\mathcal{N}$ are precisely $\{t_1, ..., t_k\}$, (ii) the map $f : N \to M$ sending a point with type t_i to m_i is a functional bisimulation. It then follows that each type t_i is isolated by its basic neighborhood $\langle\theta_i\rangle$. $\square$

As in both previous cases, let us take a model $\mathcal{M}, \mathsf{m}$ that is infinite, modally saturated, all the points are reachable from m and no two points realize different types, but the model is not bisimilar to a finite one. Again, this uses the fact that the class $\mathcal{C}_{\to^+}$ under consideration is closed under generated substructures and quotients by bisimulations. It follows that the root has infinitely many descendants. We will need the following fact:

Lemma 3.1.18. *There exists a point $\mathsf{n}_\infty \in M$ s.t. $\mathsf{n}_\infty \to^+ \mathsf{n}_\infty$ and its type t_∞ is a non-isolated element of $\{\mathbf{tp}^{\mathcal{M}}(\mathsf{n}') \mid \mathsf{n}_\infty \to^+ \mathsf{n}'\}$.*

Proof. We will inductively construct a sequence of (not necessarily distinct) points, indexed by countable ordinals $(\mathsf{n}_\alpha)_{\alpha<\omega_1} \subseteq \mathcal{M}$ with the property that for any $\alpha < \beta$: (i) $\mathsf{n}_\alpha \to^+ \mathsf{n}_\beta$ and (ii) $\mathbf{tp}^{\mathcal{M}}(\mathsf{n}_\beta)$ is not isolated in $\{\mathbf{tp}^{\mathcal{M}}(\mathsf{n}') \mid \mathsf{n}_\alpha \to^+ \mathsf{n}'\}$.

For the induction base, we simply take the root $\mathsf{n}_0 = \mathsf{m}$.

For $\alpha + 1$, we know by induction assumption that $\mathbf{tp}^{\mathcal{M}}(\mathsf{n}_\alpha)$ is non-isolated in a subset of $\mathcal{C}_{\to^+}$ and hence also non-isolated in $\mathcal{C}_{\to^+}$. Thus, by Proposition 3.1.17 we know that the model generated by n_α has to be infinite (except for the base case $\alpha = 0$ where the fact that m has infinitely many descendants is just an assumption). Now we look at the infinite set $T_\alpha = \{\mathbf{tp}^{\mathcal{M}}(\mathsf{n}')|\ \mathsf{n}_\alpha \to^+ \mathsf{n}'\}$ and pick some its limit: a non-isolated type $t_{\alpha+1} \in \mathbb{T}_{\mathcal{C}_{\to^+}}$ which, by modal saturation, is realized in some descendant $\mathsf{n}_{\alpha+1}$ of n_α.

For a limit ordinal α, we fix a subsequence $(\alpha_i)_{i\in\omega} \subseteq \alpha$ of shape ω which is cofinal with α (such subsequence exists because α is countable). Take any limit t_α of the set $T_\alpha = \{\mathbf{tp}^{\mathcal{M}}(\mathsf{n}_{\alpha_i}) \mid i \in \omega\}$. Since t_α is not isolated and $\mathbb{T}_{\mathcal{C}_{\to^+}}$ is Hausdorff, every $\varphi \in t_\alpha$ must belong to infinitely many types from T_α. It follows that there are arbitrary big i s.t. $\varphi \in t_{\alpha_i}$, so every n_{α_j}, and hence by cofinality also every n_β, has a descendant satisfying φ. Hence, by modal saturation, each n_β has a descendant realizing t_α. Moreover, since in $\mathcal{M}$ no two different points satisfy the same formulae, the point n_α realizing t_α is unique.

Now we claim that $\mathsf{n}_\alpha = \mathsf{n}_\beta$ for some $\alpha \neq \beta$. Indeed, observe that if $\mathsf{n} \to^+ \mathsf{n}'$, then n' cannot satisfy more formulae of the form $\Diamond\varphi$ than n. Since there are only countably many formulae, for sufficiently large α all $\mathbf{tp}^{\mathcal{M}}(\mathsf{n}_\alpha)$ may only differ on formulae equivalent to boolean combinations of Prop. But $\mathcal{P}(\mathsf{Prop})$ is finite, so $\mathsf{n}_\alpha = \mathsf{n}_\beta$ for some $\alpha < \beta$ and thus we put $\mathsf{n}_\infty = \mathsf{n}_\alpha$. It then follows from (i) that $\mathsf{n}_\infty \to^+ \mathsf{n}_\infty$. Finally, (ii) implies that the type t_∞ is not isolated in $\{\mathbf{tp}^{\mathcal{M}}(\mathsf{n}') \mid \mathsf{n}_\infty \to^+ \mathsf{n}'\}$, as desired. $\square$

Let us illustrate Lemma 3.1.18 with an example.

Example 3.1.19. The rightmost blue point in Example 3.1.16 was non-isolated but impossible to omit. However, the model was not modally saturated, as opposed tho the following one (again, for readability we skip some arrows and assume that the accessibility relation is the transitive closure of the depicted one):

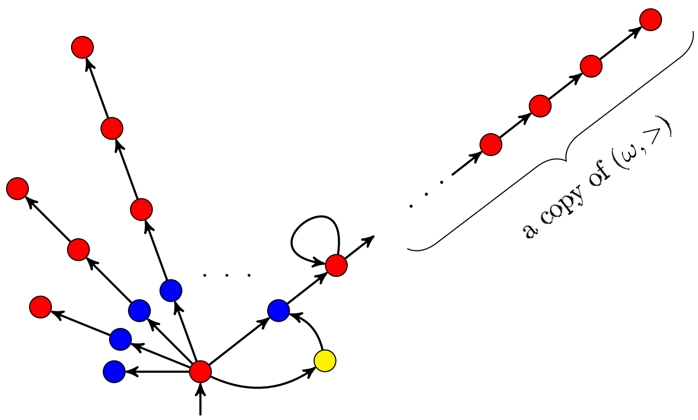

Although the type of the rightmost blue point must be realized in every model equivalent to this one, such point has infinitely many (red) children, and among these children there exists a non-isolated point (with a loop) that can be omitted.

Now, to finish the proof f Theorem 3.1.6 we define a new model $\mathcal{N}$ which has the same universe and interpretation of atomic propositions as $\mathcal{M}$ and accessibility relation obtained by removing all the arrows leading to n_∞:

$$(\to^+)^{\mathcal{N}} = (\to^+)^{\mathcal{M}} - \{(\mathsf{n}', \mathsf{n}_\infty) \mid \mathsf{n}' \in \mathcal{M}\}.$$

Observe that in $\mathcal{M}$, whenever we have $\mathsf{n} \to^+ \mathsf{n}_\infty$, then the set of descendants of n_∞ is a subset of the descendants of n and so the type t_∞ is not isolated in $\{\mathbf{tp}^{\mathcal{M}}(\mathsf{n}') \mid \mathsf{n} \to^+ \mathsf{n}'\}$. This allows us, as in the two previous cases, to prove by induction on modal depth that $\mathcal{M}, \mathsf{n} \equiv_{\mathsf{ML}} \mathcal{N}, \mathsf{n}$ for every $\mathsf{n} \in \mathcal{M}$. On the other hand, n_∞ is reachable from the root in $\mathcal{M}$ but not in $\mathcal{N}$, which gives a winning strategy for $\forall$dam in the bisimulation game. This completes the proof of Theorem 3.1.6.

3.1.3 Limitations

We end with examples illustrating limitations of our method. First of all, let us emphasize that our proofs rely on compactness of the logic under consideration and it is not hard to come up with an example of a class $\mathcal{C}$ over which ML is not compact where characterization analogous to ours fails. For instance, consider the class $\mathcal{C}_{\mathsf{WF}}$ of all well-founded models, i.e. monomodal models with no infinite paths. ML is not compact over $\mathcal{C}_{\mathsf{WF}}$ and describes the infinitely branching Hedgehog (Example 3.1.2) uniquely. This is because the only limit type, i.e. the type of any of the (bisimilar and hence equivalent) points from the limit spike, is not satisfiable in $\mathcal{C}_{\mathsf{WF}}$. In particular, this limits model theory for logics stronger than ML, such as its fixpoint extension μ-ML, which can express well-foundedness.

Since non-compact logics seem out of our reach, a natural question is if compactness is *sufficient* for an analogous characterization. Unfortunately, this is not the case. The following example shows that even the stronger assumption of first-order axiomatizability of the class $\mathcal{C}$, which implies the (relativized) van Benthem Theorem 3.1.13 and hence also compactness, is not sufficient.

Recall that in the proof of Theorem 3.1.6 we used the assumption that there are only finitely many atomic propositions. The reader could be tempted to think that with more atomic propositions it is easier to find non-isolated points, so the theorem should be true also without that assumption. Surprisingly, however, the condition turns out to be necessary.

Example 3.1.20. Assume $\mathsf{Prop} = \{\text{yellow}\} \cup \{\mathbf{blue}_\alpha \mid \alpha < \omega\}$ and consider the following model $\mathcal{M}, \mathsf{m}$:

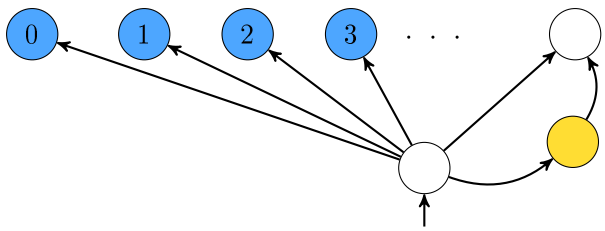

where each point satisfies at most one atomic proposition: blue point with number α satisfies $\mathbf{blue}_\alpha$, the yellow one satisfies yellow and both white ones satisfy no atomic propositions whatsoever.

We claim that every other transitive model $\mathcal{M}', \mathsf{m}'$ satisfying the same modal formulae must be bisimilar to the one above. To prove this, it suffices to come up with a winning strategy for $\exists$ve in the bisimilarity game $\mathcal{G}_{\Leftrightarrow}(\mathcal{M}, \mathcal{M}')$. This is easy and so we focus on the only interesting case when in the first move $\forall$dam chooses a point n that satisfies no atomic propositions.

First consider the case when $\mathsf{n} \in \mathcal{M}$, i.e. n is the upper white point from the picture. Since $\mathcal{M}, \mathsf{m}$, and by equivalence also $\mathcal{M}', \mathsf{m}'$, satisfy $\Diamond\Diamond\top$, $\Box\Box\Box\bot$ and $\Box\Box\neg\tau$ for every $\tau \in \mathsf{Prop}$, it follows that m' must have a grandchild n' that has no children and satisfies no atomic propositions. Hence, such n' is bisimilar to n. By transitivity, grandchildren of m' are also its children, so n' is a legal answer.

Second, consider the case with $\mathsf{n}' \in \mathcal{M}'$. Since both $\mathcal{M}, \mathsf{m}$ and $\mathcal{M}', \mathsf{m}'$ satisfy $\Box(\text{yellow} \iff \Diamond\top)$ and n' does not satisfy yellow, it does not satisfy $\Diamond\top$ either, i.e. has no children. This means that n' is bisimilar to the upper white point in $\mathcal{M}$, making it a winning answer and hence completing the proof.

The above example exploits the fact that the set of atomic propositions is infinite. A natural question is whether one can cook up a counterexample with finitely many propositions. The following example shows that even if the set Prop is *empty*, compactness is still not sufficient. Recall the class $\mathcal{C}_\forall$ of universal models. It is definable by a single first-order sentence: $\forall_{x,y} x\langle\exists\rangle y$. However, consider the following model $\mathcal{M} \in \mathcal{C}_\forall$:

Example 3.1.21. $M = \omega + 1 = \{0, 1, ..., \omega\}$; $\mathsf{m} \to^{\mathcal{M}} \mathsf{n}$ iff $\mathsf{m} = \mathsf{n} + 1$ or $\mathsf{m} = \mathsf{n} = \omega$. We assume $\mathsf{Prop} = \emptyset$.

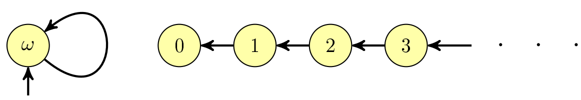

Observe that $\mathcal{M}, \mathsf{m} \not\equiv_{\mathsf{ML}} \mathcal{M}, \mathsf{n}$ for all $\mathsf{m} \neq \mathsf{n}$ and so every point has infinitely many pairwise non-equivalent children with respect to the universal modality. However, it is not hard to show that any model $\mathcal{N}, \mathsf{m}$ in $\mathcal{C}_\forall$ equivalent to $\mathcal{M}, \omega$ must be bisimilar to it. The problem is that although the topological part of our reasoning still works and we may find a limit of the types realized in $\mathcal{M}$ (in fact, in this situation there is precisely one such limit type, the type of ω) it is not possible to omit that limit type.

3.2 ML over Ordinal Models

In the previous section we have investigated bisimulational categoricity of ML over various classes of models. In each characterization, we rely on a number of good properties of ML over the class implied by its FO-axiomatizability. In particular, compactness and existence of saturated models are crucial, even if in some cases deriving the characterization requires some effort and Examples 3.1.20 and 3.1.21 demonstrate that sometimes these properties are not sufficient at all.

In this section we take different approach. We investigate the class of monomodal models where the accessibility relation is a well-founded linear order. Since every well-founded linear order is isomorphic with an ordinal, we investigate *ordinal models*. More specifically:

Definition 3.2.1. An ordinal model is a monomodal model $\mathcal{M}$ whose universe M with the only accessibility relation $\to$ are identical to some ordinal with descending (strict) order:

$$(M, \to) \;=\; (\kappa, >)$$

for some $\kappa \in \mathsf{Ord}$ called the *length* of $\mathcal{M}$ (recall that we identify every ordinal with the set of all smaller ordinals). We represent such $\mathcal{M}$ as a sequence of colors:

$$(c_\alpha)_{\alpha<\kappa} \in (\mathcal{P}(\mathsf{Prop}))^\kappa$$

and use the same additive notation as with bare ordinals: $\mathcal{M} + \mathcal{N}$ denotes the concatenation of sequences $\mathcal{M}$ and $\mathcal{N}$. The class of all ordinal models is denoted $\mathcal{C}_{\mathsf{Ord}}$.

Well-foundedness is a prototypical example of a property that is not FO-axiomatizable. Therefore, we cannot derive any desirable properties of ML from the corresponding ones for FO the way we did before. Nonetheless, in Theorem 3.2.3 we give a characterization similar to the previous ones: a theory has a unique ordinal model iff it has a finite one. Instead of using compactness of ML, the proof combines compactness of its propositional fragment with a pumping argument. An arguably more interesting (and relatively harder to prove) result is Theorem 3.2.6 characterizing compactness of ML over $\mathcal{C}_{\mathsf{Ord}}$ in terms of the number of atomic propositions: ML is compact over $\mathcal{C}_{\mathsf{Ord}}$ iff $|\mathsf{Prop}| < \omega$. Both results, apart from being interesting in their own right, demonstrate how bisimulational categoricity can be independent from compactness. On top of that, in Theorem 3.2.12 we show that ML over $\mathcal{C}_{\mathsf{Ord}}$ has a *short model property* in a suitable sense.

Bisimulation-invariant properties only depend on the descendants of the root, i.e. strictly smaller points. In order to meaningfully compare models that do not necessarily contain a greatest element, we introduce the following notation. Given models $\overline{c} = (c_\alpha)_{\alpha<\kappa}$ and $\overline{d} = (d_\alpha)_{\alpha<\lambda}$:

$$\overline{c} \sim_{\mathsf{ML}} \overline{d} \iff \overline{c} + e, \kappa \equiv_{\mathsf{ML}} \overline{d} + e, \lambda$$

where $e \in \mathcal{P}(\mathsf{Prop})$ is a fixed color. That is, we require $\overline{c}$ and $\overline{d}$ to satisfy the same modal formulae once they are both extended with a greatest element e taken as a root. Since modal equivalence of $\overline{c} + e, \kappa$ and $\overline{d} + e, \lambda$ does not depend on the specific choice of e, we leave it unspecified. It follows that the two conditions are equivalent:

- $\overline{c} \sim_{\mathsf{ML}} \overline{d}$;
- for every $k < \omega$ and finite $P \subseteq \mathsf{Prop}$:

$$\overline{c} + e, \kappa \leftrightarroweq^{k,P} \overline{d} + e, \lambda.$$

In the light of (2.8), we usually think of the relation $\leftrightarroweq^{k,P}$ in terms of the restricted k-step bisimilarity game $\mathcal{G}_{\leftrightarroweq}^{<\omega,P}(\overline{c}+e, \overline{d}+e), (\kappa, \lambda, k, \mathsf{count})$. Recall that $\mathcal{G}_{\leftrightarroweq}^{<\omega,P}$ is the same game as $\mathcal{G}_{\leftrightarroweq}^{<\omega}$ except that only propositions from P are checked in the base stage. Let us use this equivalence to prove that $\sim_{\mathsf{ML}}$ is a congruence for concatenation:

Proposition 3.2.2. *Assume two sequences $\overline{c}_0, \overline{c}_1, \ldots$ and $\overline{d}_0, \overline{d}_1, \ldots$ of models, both of length $\xi \in \mathsf{Ord}$, satisfying $\overline{c}_\gamma \sim_{\mathsf{ML}} \overline{d}_\gamma$ for every $\gamma < \xi$. Then:*

$$\overline{c} = \sum_{\gamma\in[0,\xi)} \overline{c}_\gamma, \sim_{\mathsf{ML}} \sum_{\gamma\in[0,\xi)} \overline{d}_\gamma = \overline{d}.$$

Proof. Fix finite P. For every $k < \omega$ and $\alpha < |\overline{c}_\gamma|$, the assumption $\overline{c}_\gamma \sim_{\mathsf{ML}} \overline{d}_\gamma$ implies existence of $\beta < |\overline{d}_\gamma|$ such that $\overline{c}_\gamma, \alpha \leftrightarroweq^{k,P} \overline{d}_\gamma, \beta$. Symmetrically, for every $\beta < |\overline{d}_\gamma|$ and $k < \omega$ there is $\alpha < |\overline{c}_\gamma|$ with the same property. Denote the lengths of $\overline{c}$ and $\overline{d}$ by κ and λ. We show that $\overline{c} + e, \kappa \leftrightarroweq^{k,P} \overline{d}, \lambda$ for all $k < \omega$ by giving $\exists$ve a winning strategy from $(\kappa, \lambda, k, \mathsf{count})$. The strategy is easy: preserve the invariant that both points α and β come from $\overline{c}_\gamma$ and $\overline{d}_{\gamma'}$ for the same $\gamma = \gamma'$ and $\overline{c}_\gamma, \alpha \leftrightarroweq^{k,P} \overline{d}_\gamma, \beta$. ☐

3.2.1 Bisimulational Categoricity

Let us start with an analysis of bisimulational categoricity over $\mathcal{C}_{\mathsf{Ord}}$.

Theorem 3.2.3. *For every type $t \in \mathbb{T}_{\mathcal{C}_{\mathsf{Ord}}}$, the following are equivalent:*

(1) *t has a unique ordinal model up to $\leftrightarroweq$;*

(2) *every ordinal model of t is bisimilar to a finite ordinal model;*

(3) *t has a finite ordinal model.*

Note that given pointed models $\overline{c}, \kappa$ and $\overline{d}, \lambda$, a straightforward transfinite induction on $\mathsf{max}(\kappa, \lambda)$ reveals that $\overline{c}, \kappa$ and $\overline{d}, \lambda$ are bisimilar iff their parts $(c_\alpha)_{\alpha \in [0,\kappa)}$ and $(d_\beta)_{\beta \in [0,\lambda)}$ consisting of points below the root are identical. Thus, if we require a model to be reachable from the root (i.e. require the root to be the greatest element), which we can do without losing generality, then we would get that t has a unique ordinal model iff all its ordinal models are finite iff it has a finite ordinal model. Let us prove the theorem.

Proof. As before, the implication (2) $\implies$ (3) is obvious and (3) $\implies$ (1) follows from Proposition 3.1.7 so it remains to prove (1) $\implies$ (2). Assume that t has a model $\overline{c}, \kappa$ that is not bisimilar to a finite ordinal model, meaning that κ is infinite. It suffices to construct a model for t not bisimilar to $\overline{c}, \kappa$. This is an immediate corollary of the following proposition.

Proposition 3.2.4 (Upper Skolem-Löwenheim for $\mathcal{C}_{\mathsf{Ord}}$). *Assume a modal type t. If t is satisfiable in an ordinal model $\overline{c}, \kappa$ with infinite κ then it has ordinal models $\overline{d}, \lambda$ for arbitrarily big λ.*

Proof. In the case of ML (or any bisimulation-invariant formalism) over a class of models $\mathcal{C}$, such result in the style of the upper Skolem-Löwenheim Theorem is often immediate: if $\mathcal{C}$ is closed under (infinite) disjoint unions the result is trivial; if $\mathcal{C}$ is FO-axiomatizable then by Theorem 3.1.13 it follows from an analogous result for FO. The class $\mathcal{C}_{\mathsf{Ord}}$ is neither of the two cases, however, so we need a separate proof. Assume $\overline{c}, \kappa \models t$ as in the formulation. Without loss of generality the root is the greatest element, i.e. $\overline{c} = (c_\alpha)_{\alpha \leq \kappa}$.

Identifying every color $c \in \mathcal{P}(\mathsf{Prop})$ with its characteristic function $\chi_c : \mathsf{Prop} \to \{0,1\}$ we may view the set of all colors as a compact topological space $\{0,1\}^{\mathsf{Prop}}$. Hence, there exists a color $c_{\mathrm{lim}} \in \mathcal{P}(\mathsf{Prop})$ that is a limit of the colors $(c_\alpha)_{\alpha<\omega}$. Explicitly, this means that for every finite $P \subseteq \mathsf{Prop}$ there are arbitrarily big $\alpha < \omega$ such that $\chi_{c_{\mathrm{lim}}}$ and χ_{c_α} are equal on P.

For every ordinal $\xi \in \mathsf{Ord}$, consider the model:

$$(c_{\mathrm{lim}})^\xi = c_{\mathrm{lim}} c_{\mathrm{lim}} \cdots$$

being $(\xi, >)$ with every point colored by c_{lim}. We claim that inserting ξ-many copies of c_{lim} between $[0, \omega)$ and $[\omega, \kappa]$:

$$(c_\alpha)_{\alpha \in [0,\omega)} + (c_{\mathrm{lim}})^\xi + (c_\alpha)_{\alpha \in [\omega,\kappa]} = (d_\beta)_{\beta \in [0,\lambda]} = \overline{d}$$

results in a modally equivalent model:

$$\bar{c}, \kappa \equiv_{\mathsf{ML}} \bar{d}, \lambda \tag{3.1}$$

regardless of the choice of ξ. The above λ is the (unique) ordinal isomorphic to the concatenation of orders ω, ξ and $[\omega, \kappa]$. Thus, taking ξ of arbitrarily big cardinality we get arbitrarily big λ and infer Proposition 3.2.4 (and therefore also Theorem 3.2.3) from (3.1).

The equivalence follows from $(c_\alpha)_{\alpha\in[0,\kappa)} \sim_{\mathsf{ML}} (d_\beta)_{\beta\in[0,\lambda)}$ because the colors c_κ and d_λ are equal. Due to Proposition 3.2.2 it suffices to show:

$$(c_\alpha)_{\alpha\in[0,\omega)} \sim_{\mathsf{ML}} (c_\alpha)_{\alpha\in[0,\omega)} + (c_{\mathrm{lim}})^{\xi}.$$

To that end, fix finite $P \subseteq \mathsf{Prop}$ and $k < \omega$. We prove:

$$(c_\alpha)_{\alpha\in[0,\omega)} + e \leftrightarrows^{k,P} (c_\alpha)_{\alpha\in[0,\omega)} + (c_{\mathrm{lim}})^{\xi} + e$$

by giving $\exists$ve a strategy winning from $(\omega, \omega{+}\xi, k, \mathsf{count})$ in the corresponding game.

For each $\alpha < \omega{+}\xi$, denote by $c^{\circ}_{\alpha} = \pi(c_\alpha)$ the image under the projection $\pi : \mathcal{P}(\mathsf{Prop}) \to \mathcal{P}(P)$. The set $\mathcal{P}(P)$ of colors is finite. Let $x_0 < \omega$ be the least number such that all the colors appearing finitely often in $(c^{\circ}_{\alpha})_{\alpha\in[0,\omega)}$ appear below x_0 and for $i < \omega$ let x_{i+1} be the least $x < \omega$ such that all colors appearing infinitely often in $(c^{\circ}_{\alpha})_{\alpha\in[0,\omega]}$ are present in $(x_i, x]$.

This leads to the following strategy for $\exists$ve. She preserves the invariant that if the current pair of points is (α, β) and there are $k < \omega$ rounds left before the end of the game, then either (i) $\alpha = \beta$ or (ii) the points satisfy the same atomic propositions from P, and $x_l < \alpha, \beta$. This is always possible because by definition of c_{lim}, its projection $\pi(c_{\mathrm{lim}}) = c^{\circ}_{\mathrm{lim}}$ appears infinitely often in $(c^{\circ}_{\alpha})_{\alpha\in[0,\omega)}$. □

□

3.2.2 Compactness

Let us now analyze compactness of ML over $\mathcal{C}_{\mathsf{Ord}}$. It is easy to see that if the set Prop is infinite then the logic is not compact.

Example 3.2.5. Assume $\mathsf{Prop} = \{\tau_i \mid i < \omega\}$ and consider the set of modal sentences:

$$t = \{\Box(\tau_i \implies \Diamond\tau_{i+1}) \mid i < \omega\} \cup \{\Diamond\tau_0\}.$$

That is, for each i we require that if a point α_i satisfies τ_i then some $\alpha_{i+1} < \alpha_i$ satisfies τ_{i+1}, and some point satisfies the first τ_0. It follows that t cannot be satisfied in a well-founded model because it would imply existence of an infinite descending chain $\alpha_0 > \alpha_1 > \ldots$ with each α_i satisfying τ_i.

On the other hand, any finite subset $t_0 \subseteq t$ is satisfiable in $\mathcal{C}_{\mathsf{Ord}}$. Such finite t_0 uses only finitely many propositions. Let $l < \omega$ be the maximal i for which τ_i appears in t_0. Consider the finite ordinal model of shape $(\{l+1, ..., 0\}, >)$ with $i \models \tau_j$ iff $l - i = j$. Such model, with $l+1$ as a root, satisfies t_0. It follows that every finite fragment of t has a model in $\mathcal{C}_{\mathsf{Ord}}$ but the entire t does not.

Infiniteness of Prop is crucial in the above example. In the case of FO, this assumption can be easily dropped: even with $\mathsf{Prop} = \emptyset$ (which means that ordinal models $\mathcal{C}_{\mathsf{Ord}} = \mathsf{Ord}$ are just ordinals) there exists a set of FO-sentences whose every fragment is satisfiable in Ord but the entire set is not. FO can express finiteness of the model by saying that there is exactly one point with no predecessor. Then, a sequence of sentences $\varphi_1, \varphi_2, ...$, with each φ_i stating that there are at least i distinct elements, witnesses the lack of compactness.

Somehow surprisingly, however, it turns out that in the case of ML if the set Prop is finite then the logic is compact over $\mathcal{C}_{\mathsf{Ord}}$.

Theorem 3.2.6. *Assume that* $|\mathsf{Prop}| < \omega$ *and take a set of modal sentences* t*. If every finite fragment of* t *is satisfiable in* $\mathcal{C}_{\mathsf{Ord}}$ *then so is the entire* t*.*

Towards the proof, let us introduce some terminology. Consider the set:

$$t_{\mathsf{WF}} = \{\Diamond\varphi \implies \Diamond(\varphi \wedge \Box\neg\varphi) \mid \varphi \in \mathsf{ML}\}$$

of modal sentences. The meaning of this is that whenever a point has a descendant α satisfying φ then there is a *least* $\beta \leq \alpha$ satisfying it. Since every subset of an ordinal model (in particular the set $[\![\varphi]\!]$ of points satisfying φ) has the least element, every point of a model in $\mathcal{C}_{\mathsf{Ord}}$ satisfies t_{WF}.

A *proto-model* is a monomodal model $\mathcal{M}$ which is almost an ordinal model except that the assumption of well-foundedness is weakened and we only require t_{WF} to be valid. That is, the accessibility relation of a proto-model equals $>$ for a linear order $\leq$ on the universe and each point satisfies t_{WF}. Since proto-models are closed under appending a single greatest element, the notation $\sim_{\mathsf{ML}}$ extends from models to proto-models in an obvious way. We call a point m in a (proto-)model $\mathcal{M}$ *modally definable* if there is a modal formula φ defining it in the sense that m is the unique point in $\mathcal{M}$ satisfying φ: $[\![\varphi]\!] = \{\mathsf{m}\}$. We call $\mathcal{M}$ definable if all its points are definable. Let us prove the following lemma.

Lemma 3.2.7. *Given a proto-model* $\mathcal{M}$*, let* $\mathcal{M}^{def}$ *be its submodel with universe* M^{def} *consisting of modally definable points of* $\mathcal{M}$ *and the order and colors inherited from* $\mathcal{M}$*. We have:*

$$\mathcal{M}, \mathsf{m} \equiv_{\mathsf{ML}} \mathcal{M}^{def}, \mathsf{m}$$

for every $\mathsf{m} \in M^{def}$*. In particular,* $\mathcal{M}^{def}$ *is a definable proto-model and* $\mathcal{M} \sim_{\mathsf{ML}} \mathcal{M}^{def}$*.*

Proof. The proof of the equivalence proceeds by induction on the complexity of formulae. The only nontrivial case is to show that if $\mathcal{M}, \mathsf{m} \models \Diamond\varphi$ then also $\mathcal{M}^{\mathrm{def}}, \mathsf{m} \models \Diamond\varphi$, as in $\mathcal{M}^{\mathrm{def}}$ the point m can have fewer descendants than in $\mathcal{M}$. However, since $\mathcal{M}$ is a proto-model we have $\mathcal{M}, \mathsf{m} \models \Diamond\varphi \implies \Diamond(\varphi \wedge \Box\neg\varphi)$ and thus also $\mathcal{M}, \mathsf{m} \models \Diamond(\varphi \wedge \Box\neg\varphi)$. It follows that there must be $\mathsf{n} < \mathsf{m}$ in $\mathcal{M}$ which is the least point of $\mathcal{M}$ satisfying φ. This means that n is definable by the sentence $\varphi \wedge \Box\neg\varphi$ and so it belongs to M^{def}. Since by the induction hypothesis $\mathcal{M}^{\mathrm{def}}, \mathsf{n} \models \varphi$, we get $\mathcal{M}^{\mathrm{def}}, \mathsf{m} \models \Diamond\varphi$. $\Box$

The next proposition is the key to Theorem 3.2.6.

Proposition 3.2.8. *Assume that* $\mathsf{Prop} < \omega$. *Then every definable proto-model is an ordinal model (i.e. is well-founded).*

Let us first show how the theorem follows from Lemma 3.2.7 and Proposition 3.2.8.

Proof. Fix finite Prop and assume a set of modal sentences t whose every finite fragment is satisfiable in $\mathcal{C}_{\mathsf{Ord}}$. Being a proto-model is axiomatizable by a set of FO-sentences $T_{\mathsf{WF+lin}} \subseteq \mathsf{FO}$: the fact that $\leq$ is a linear order can be expressed by a sentence φ_{lin} and validity of t_{WF} by the set of sentences $\{\forall_x.\mathsf{ST}_x(\varphi) \mid \varphi \in t_{\mathsf{WF}}\}$. Denote by T the standard translation $\mathsf{ST}_x[t] \subseteq \mathsf{FO}$ of t. Since every finite fragment of t is satisfiable in $\mathcal{C}_{\mathsf{Ord}}$, so is every finite fragment of $T_{\mathsf{WF+lin}} \cup T$. By compactness of FO, this gives us a model (in the general sense, not necessarily belonging to $\mathcal{C}_{\mathsf{Ord}}$) $\mathcal{M}$ satisfying $T_{\mathsf{WF+lin}} \cup T$ with the variable x interpreted as some $\mathsf{m} \in \mathcal{M}$. This means that $\mathcal{M}, \mathsf{m}$ is a proto-model satisfying t. Without losing generality we assume that m is the greatest element of $\mathcal{M}$ (otherwise take the induced sub-proto-model consisting of points smaller than m).

Let $\mathcal{M}_0$ be the sub-proto-model of $\mathcal{M}$ with m removed (i.e. $M_0 = M - \{\mathsf{m}\}$). Using Lemma 3.2.7 we may take the definable sub-proto-model $\mathcal{M}_0^{\mathrm{def}}$ of $\mathcal{M}_0$ so that $\mathcal{M}_0 \sim_{\mathsf{ML}} \mathcal{M}_0^{\mathrm{def}}$. In particular, $\mathcal{M}', \mathsf{m} \models t$ where $\mathcal{M}'$ is the proto-model obtained from $\mathcal{M}_0^{\mathrm{def}}$ by appending back the point m as the greatest element. By Proposition 3.2.8, $\mathcal{M}_0^{\mathrm{def}}$ is well-founded and hence so is $\mathcal{M}'$. This means that $\mathcal{M}', \mathsf{m}$ is an ordinal model satisfying t, as desired. $\Box$

It remains to prove Proposition 3.2.8. We start by showing how any formula φ satisfied by a point m in a proto-model $\mathcal{M}$ can be rewritten into a *linear* φ'. The resulting φ' wll be potentially stronger than the original φ but still satisfied by m. The purpose of this is to replace arbitrary modal definitions with linear ones, although the translation is valid for arbitrary formulae, not necessarily definitions. For convenience, we work with the alternative syntax of ML with colors in place of of literals. Since the set Prop is finite, such syntax is equivalent to the original one.

We call a modal formula φ *linear* if there is a descending chain of subformulae $\psi_1, ..., \psi_l$ with $\varphi = \psi_1$ and:

$$\begin{aligned} \psi_i =& c_i \wedge \Box\theta_i \wedge \Diamond\psi_{i+1} && \text{for } 1 \leq i < l, \\ \psi_l =& c_l \wedge \Box\theta_l; \end{aligned}$$

where, for each $i \leq l$, c_i is a color and θ_i is an arbitrary formula. The key property of such linear formula is that the set $\{\psi_2, ..., \psi_l\}$ of all immediate subformulae of a $\Diamond$ which are not subformulae of a $\Box$ is linearly ordered by the relation of being a subformula.

Satisfaction of linear formulae has a particularly simple description. Consider a linear φ with subformulae $\psi_1, ..., \psi_l$ witnessing its linearity. For every m, $\mathcal{M}, \mathsf{m} \models \varphi$ is equivalent to the existence of a map $h : \{1, ..., l\} \to M$ such that:

1. $h(i)$ satisfies c_i and $\Box\theta_i$ for each $i \leq l$;
2. h is strictly antitone meaning that $i < j$ implies $h(i) > h(j)$ and
3. $h(1) = \mathsf{m}$.

Let us make formulae linear.

Proposition 3.2.9. *If a point* m *in a proto-model* $\mathcal{M}$ *satisfies a formula* φ *then there is a linear* φ' *satisfied by* m *and stronger than* φ *(meaning that every pointed proto-model satisfying* φ' *satisfies* φ*).*

Proof. Assume a pointed proto-model $\mathcal{M}, \mathsf{m}$ and a formula φ true at m.

We first show how to remove all the disjunctions that do not appear in the scope of at least one $\Box$ operator. If no superformula of $\psi \vee \psi' \in \mathsf{SubFor}(\varphi)$ begins with $\Box$ then we obtain a stronger formula that is still satisfied in m by appropriately choosing one of the disjuncts, ψ or ψ', and putting it in place of $\psi \vee \psi'$. The intuition is that if the only modal operators above $\psi \vee \psi'$ are $\Diamond$'s but not $\Box$'s then it suffices to have one witness n for $\psi \vee \psi'$ in the model; and if n satisfies $\psi \vee \psi'$ then it satisfies at least one of the disjuncts.

Formally, consider the semantic game corresponding to the evaluation of φ from m. By the Adequacy Theorem 2.3.1 for ML, $\exists$ve has a winning strategy σ in that game. Since no $\Box$ appears in the (unique) path from φ to $\psi \vee \psi'$, whenever a σ-play π reaches a position of shape $(\mathsf{n}, \psi \vee \psi')$ then all the $\forall$dam's choices in π are propositional and thus they are all encoded in the subformula $\psi \vee \psi'$ itself. Consequently, there is at most one σ-play π ending in $\psi \vee \psi'$. The strategy wins by picking either ψ or ψ' after that π. Hence, if we remove the other disjunct the same σ will witnesses that $\exists$ve wins the semantic game for the resulting formula. We call such removal of not-boxed disjunctions *dedisjunctification*.

Consider two rewriting rules:

$$\begin{array}{lcl} \Box\chi \wedge \Box\chi' & \mapsto & \Box(\chi \wedge \chi') \\ \Diamond\chi \wedge \Diamond\chi' & \mapsto & \Diamond(\chi \wedge \chi') \vee \Diamond(\chi \wedge \Diamond\chi') \vee \Diamond(\Diamond\chi \wedge \chi'). \end{array}$$

Both rules produce semantically equivalent formulae. The first one is valid over arbitrary models. The second one exploits linearity of the accessibility relation $<$: if m has descendants n satisfying ψ and n' satisfying ψ' then these descendants are either equal or ordered and the three disjuncts correspond to the three possible cases $\mathsf{n} = \mathsf{n}'$, $\mathsf{n}' < \mathsf{n}$ and $\mathsf{n} < \mathsf{n}'$.

Using the two above rules and dedisjunctification we turn φ into the desired linear form. Starting from $k = 1$, after the k-th step we want to have $\psi_1, ..., \psi_k$ with $\varphi = \psi_1$ and $\psi_i = c_i \wedge \Box\theta_i \wedge \Diamond\psi_{i+1}$ for all $1 \leq i < k$. That is, φ is almost linear except that the last ψ_k can be arbitrary. Initially, the condition is trivial with $k = 1$ and just one $\psi_1 = \varphi$.

Assume that after the k-th step we have the mentioned $\psi_1, ..., \psi_k$. It suffices to rewrite the deepest ψ_k into $c_k \wedge \Box\theta_k \wedge \Diamond\psi_{k+1}$ with arbitrary ψ_{k+1}. By dedisjunctification, we may remove all the disjunctions from ψ_k so that it becomes a conjunction $\chi_1 \wedge ... \wedge \chi_d$ of colors and formulae beginning with $\Diamond$ and $\Box$. Without loss of generality, at least one conjunct begins with a $\Box$ (otherwise we can add $\Box\top$ which is equivalent to $\top$) and precisely one color (conjunction of at least two different colors is inconsistent and if there are no colors then we could take a disjunction of all the colors and apply dedisjunctification to pick precisely one of them). Hence, we get a formula of the form $c_k \wedge \Box\chi_1^{\Box} \wedge ... \wedge \Box\chi_m^{\Box} \wedge \Diamond\chi_1^{\Diamond} \wedge ... \wedge \Diamond\chi_n^{\Diamond}$. Using the first rule we replace all $\Box\chi_1^{\Box} \wedge ... \wedge \Box\chi_m^{\Box}$ with a single $\Box(\chi_1^{\Box} \wedge ... \wedge \chi_m^{\Box}) = \Box\theta_{k+1}$. If no conjunct begins with $\Diamond$ the entire procedure terminates and if there is exactly one $\Diamond\chi_1^{\Diamond}$ we are done with the $k+1$-st step putting $\psi_{k+1} = \chi_1^{\Diamond}$. Otherwise, for every pair of conjuncts beginning with $\Diamond$ we (i) use the second rule to replace it with a disjunction $\Diamond\xi_1 \vee \Diamond\xi_2 \vee \Diamond\xi_3$ of three formulae beginning with $\Diamond$ and immediately (ii) dedisjunctify to get just one $\Diamond\xi_1, \Diamond\xi_2$ or $\Diamond\xi_3$.

It remains to prove that the rewriting terminates. The second rewriting rule can increase the nesting of the $\Diamond$ operators. However, since every its application is followed by dedisjunctification, the overall number of $\Diamond$ operators in the formula never increases and after the $k+1$-st step the last formula ψ_{k+1} has strictly less $\Diamond$'s than ψ_k after k-th step. Thus, the procedure terminates after at most $|\varphi|$ steps, bringing φ into the desired linear form. $\Box$

We finish the proof of Proposition 3.2.8 and hence also Theorem 3.2.6.

Proof. Assume a definable proto-model $\mathcal{M}$ over Prop. Each point m in $\mathcal{M}$ comes with a modal formula φ defining it. Thanks to Proposition 3.2.9 we assume that all these definitions are linear.

Assume towards contradiction that $\mathcal{M}$ is not well-founded: there exists an infinite descending chain $\mathsf{m}_1 > \mathsf{m}_2 > ...$ in $\mathcal{M}$. For each m_i with a linear definition φ presented as $\psi_1, ..., \psi_l$ let $w_i = c_1...c_l \in (\mathcal{P}(\mathsf{Prop}))^*$ be the consecutive colors from the sequence. This way, we get an infinite sequence of finite words $w_1, w_2, ...$ over $\Gamma = \mathcal{P}(\mathsf{Prop})$. Consider the binary *subsequence* relation $\preceq_{\mathsf{ss}}$ on Γ^* defined as $d_1...d_m \preceq_{\mathsf{ss}} e_1...e_n$ iff there is a subset $i_1 < ... < i_m$ of positions $1, ..., n$ such that $d_1...d_m = e_{i_1}...e_{i_m}$. By the famous result known as Higman's Lemma [17, Theorem 4.3], Γ^* equipped with $\preceq_{\mathsf{ss}}$ is a *well-quasi-order* meaning that every infinite sequence $v_1, v_2, ...$ in Γ^* contains a pair $v = v_i$ and $v' = v_j$ such that that $v \preceq_{\mathsf{ss}} v'$ and $i < j$. In particular, our sequence $w_1, w_2, ...$ contains such a pair. This means that there are points $\mathsf{m} > \mathsf{m}'$ in the sequence with linear definitions φ and φ' whose corresponding words w and w' satisfy $w \preceq_{\mathsf{ss}} w'$. Let $h : \{1, ..., l\} \to M$ and $h' : \{1, ..., l'\} \to M$ be the antitone maps witnessing that the two points satisfy φ and φ'. By $w \preceq_{\mathsf{ss}} w'$, there exists a subsequence $i_1 < ... < i_l$ of $1, ..., l'$ such that $c_j = c'_{i_j}$ for all $j \leq l$. Consider the map $h'' : \{1, ...l\} \to M$ given by:

$$h''(j) = \mathsf{min}(h(j), h'(i_j))$$

for every j. By monotonicity of min, such h'' is strictly antitone. For all $j \leq l$, $h(j)$ and $h'(i_j)$, and therefore also $h''(j)$, have the same color c_j. Moreover, the meaning of any formula beginning with a $\Box$ is a downward closed subset of M, so $h(j) \models \Box\theta_j$ and $h''(j) \leq h(j)$ imply $h''(j) \models \Box\theta_j$. It follows that the map h'' is a witness for $h''(1) \models \varphi$. However:

$$h''(1) \leq h'(i_1) \leq h'(1) = \mathsf{m}' < \mathsf{m}$$

contradicting the assumption that m is the unique point satisfying φ. $\Box$

3.2.3 Short Model Property

In the remaining part of this section we investigate what we call a *short model* property: if a modal theory t has an ordinal model then it has one that is not very long. Since well-foundedness is preserved under taking submodels, the definable submodel of any given model from $\mathcal{C}_{\mathsf{Ord}}$ belongs to $\mathcal{C}_{\mathsf{Ord}}$ as well. Given any ordinal model we can always take away the root, pick a definable submodel of the remaining part and plug the root back. Hence, Lemma 3.2.7 implies that any t satisfiable in $\mathcal{C}_{\mathsf{Ord}}$ is satisfied in a model where every point is definable, with a possible exception for the root. Since the size of a definable model is not greater than the number of formulae we get the following proposition, complementary to Proposition 3.2.4, as a corollary.

Proposition 3.2.10 (Lower Skolem-Löwenheim for $\mathcal{C}_{\mathsf{Ord}}$)**.** *Assume a modal theory t satisfiable in $\mathcal{C}_{\mathsf{Ord}}$. The theory has an ordinal model of cardinality not greater than* $\mathsf{max}(\omega, |\mathsf{Prop}|)$.

Ordinal models come equipped with ordinal-valued length, a more fine-grained measure than the cardinal-valued size. With infinite Prop, the above proposition is already optimal for such length: for arbitrary ordinal $\kappa \in \mathsf{Ord}$ there is an ordinal model of length $\kappa + 1$ over $|\mathsf{Prop}| = |\kappa|$ with no shorter modally equivalent ordinal model. Take the set $\mathsf{Prop} = \{\tau_\alpha \mid \alpha < \kappa\}$ and an ordinal model $\overline{c}_\kappa, \kappa$ of shape $\kappa + 1$ where a point α satisfies τ_β iff $\alpha = \beta$. It is easy to see that any model modally equivalent to $\overline{c}_\kappa, \kappa$ has it as a submodel.

With finite Prop more can be said. For a more precise analysis, for the rest of this section we will refer to the number of colors instead of the number of propositions. With k atomic propositions there are 2^k colors, but it is insightful to look at models with l colors for l that is not necessarily a power of 2. The following example gives a lower bound on the length of models.

Example 3.2.11. For every $1 \leq k < \omega$ consider colors $c_1, c_2 ....$ We inductively define sequences $w_0, w_1, ...$:

$$\begin{aligned} w_0 &= \varepsilon \\ w_{k+1} &= (w_k + c_{k+1})^\omega. \end{aligned}$$

For each $k < \omega$, w_k uses colors $c_1, ..., c_k$, has length ω^k and there is no sequence $\overline{d}$ shorter than that with $w_k \sim_{\mathsf{ML}} \overline{d}$. The last property follows from the invariant that every point α in each sequence w_k is modally definable. Such invariant implies the property because in any $\overline{d} \sim_{\mathsf{ML}} w_k$ the points β and β' satisfying the respective definitions φ and φ' of α and α' from w_k must have the same order as α and α'.

An arguably more interesting part is the following upper bound matching the above lower one.

Theorem 3.2.12 (Short Model Property)**.** *Assume a modal theory t over $k < \omega$ colors satisfiable in $\mathcal{C}_{\mathsf{Ord}}$. The theory has an ordinal model of length at most $\omega^k + 1$.*

Proof. It suffices to show that for every $\overline{c}$ over k colors there is $\overline{d}$ of length at most ω^k with $\overline{c} \sim_{\mathsf{ML}} \overline{d}$. We prove that by induction. The base step with $k = 0$ is trivial, as the only sequence over the empty set of colors is the empty one ε. Assume the claim is true for k and take a model $\overline{c} = (c_\alpha)_{\alpha<\kappa}$ over $k + 1$ colors. For every $i < \omega$ inductively define:

$$\begin{aligned} \alpha_0 &= 0 \\ \alpha_{i+1} &= \text{(if it exists) the least } \alpha \text{ such that all the colors appear in } (\alpha_i, \alpha]. \end{aligned}$$

Let z denote either ω if there are infinitely many α_i's or $i + 1$ if α_i is the greatest one. Put:

$$\bigcup_{i<z} [0, \alpha_i) = (c_\alpha)_{\alpha<\kappa'} = \overline{c}'.$$

We claim that:

$$\overline{c} \sim_{\mathsf{ML}} \overline{c}'. \tag{3.2}$$

This is nontrivial only with $z = \omega$ as otherwise $\overline{c} = \overline{c}'$. If $z = \omega$ we prove the equivalence by showing that $\overline{d} + e \leftrightarroweq^n \overline{d} + e$ for all $n < \omega$. To that end, we give $\exists$ve a winning strategy in the game $\mathcal{G}_{\leftrightarroweq}^{<\omega}(\overline{c}, \overline{c}')$ from $(\kappa, \kappa', n, \mathsf{count})$. The strategy is easy: preserve the invariant that if the current points are (α, α') and there are l rounds left then either $\alpha = \alpha'$ or the points have the same color and $\alpha_l < \alpha, \alpha'$.

Denote $I_i = (\alpha_i, \alpha_{i+1})$ if $i+1 < z$ and $I_i = (\alpha_i, \kappa]$ if $i+1 = z$. Each $\overline{c}_i = (c_\alpha)_{\alpha \in I_i}$ contains at most k colors. Thus, by the induction hypothesis there exists $\overline{d}_i$ of length $\kappa_i \leq \omega^k$ such that $\overline{c}_i \sim_{\mathsf{ML}} \overline{d}_i$. Applying Proposition 3.2.2 we get:

$$\overline{c}' = \sum_{i<z}(c_{\alpha_i} + \overline{c}_i) \sim_{\mathsf{ML}} \sum_{i<z}(c_{\alpha_i} + \overline{d}_i) = \overline{d}$$

which, together with (3.2), leads to:

$$\overline{c} \sim_{\mathsf{ML}} \overline{d}.$$

Since the length of $\overline{d}$ is:

$$\sum_{i<z}(1 + \kappa_i) \leq \sum_{i<z}(1 + \omega^k) = \sum_{i<z} \omega^k \leq \omega^{k+1},$$

this completes the proof. $\square$

Chapter 4

Countdown μ-calculus

The modal μ-calculus has a number of desirable properties such as simple syntax and a tight (and effective) connection with simple parity games and automata, which make it a convenient formalism to study. At the same time, the logic is rather expressive: it can define all bisimulation-invariant properties definable in monadic second-order logic (MSO) [20, Theorem 11], such as "there is an infinite path of a-labeled edges". However, there are some properties of interest which are not definable even in MSO. Notable examples include (un)boundedness properties such as "for every number n, there is a path with at least n consecutive a-labeled edges". An extension of MSO called MSO+U, aimed at defining such properties, has been considered [8]. However, the satisfiability problem of MSO+U turned out to be undecidable even for word models [6, Theorem 1.1]. Since the modal μ-calculus is a fragment of MSO, it is worthwhile to extend it with a mechanism for defining (un)boundedness properties, in the hope of retaining decidability.

Countdown μ-calculus. In this chapter we investigate such an extension: the *countdown μ-calculus* $\mu^{<\infty}$-ML. In addition to μ-calculus operators, it features countdown operators μ^α and ν^α parametrized by ordinal numbers α. Instead of least and greatest fixpoints, they define ordinal approximations of those fixpoints. Intuitively, while the meaning of classical μ-calculus formulae $\mu x.\varphi(x)$ and $\nu x.\varphi(x)$ is defined by infinite unfolding of the formula φ until a fixpoint is reached, for $\mu^\alpha x.\varphi(x)$ and $\nu^\alpha x.\varphi(x)$ the unfolding stops after α steps (which makes a difference if α is smaller than the closure ordinal of φ). The classical fixpoint operators are kept but renamed to μ^∞ and ν^∞, to make clear the lack of any restrictions on the unfolding process.

Countdown Games. An inductive definition of the semantics of countdown formulae is just as straightforward as in the classical case. With some more effort games are adapted to such countdown setting as well. We present *countdown games* which are similar to simple parity games known from the classical setting, but are additionally equipped with counters that are decremented and reset by the two players according to specific rules. We first describe games corresponding to fixpoint approximations F_μ^α and F_ν^α with α smaller than ∞. The new games for F_μ^α and F_ν^α extend the ones for the fixpoints F_μ^∞ and F_ν^∞ with a single counter. We illustrate the modification with the full bisimilarity $\leftrightarroweq$ and its depth-k variant $\leftrightarroweq^k$. The same way as the game characterizing $\leftrightarroweq$ is derived as an instance of a fixpoint game, in Example 4.1.3 we derive the game for $\leftrightarroweq^k$ as an instance of our fixpoint-approximation game.

The general countdown games are introduced as a nested version of these for fixpoint approximations F_μ^α and F_ν^α, similarly to simple parity games being a nested version of games for the fixpoints F_μ^∞ and F_ν^∞.

Countdown Automata. Countdown games give raise to the notion of *countdown automata*. Countdown automata extend the usual parity automata with counters: the semantic games they induce are countdown games rather simple parity ones. The correspondence between countdown formulae and such countdown automata and games is as tight as for the classical μ-calculus. Effective language-preserving translations from logic to automata and back are described in the respective Subsections 4.4.1 and 4.4.2.

Vectorial vs Scalar Calculus. The threefold correspondence between logic, games and automata lifts to the countdown setting. However, complications arise: the distinction between *vectorial* and *scalar* formulae, which in the classical case disappears to a large extent due to the Bekić principle (2.9), now becomes pronounced. While it is relatively easy to come up with a counterexample to a countdown version of the Bekić principle, this does not exclude possibility of another completely different translation. In Theorem 4.6.2 we show that vectorial countdown calculus is indeed more expressive than its scalar fragment.

Stacked Counters. We further introduce automata with stacked counters. With countdown automata, all the modifications to the counter values are inherently entangled with the visited ranks. *Automata with stacked counters* are an alternative automata model where the counters can be manipulated according to explicit instructions given by the transition function. Unlike with countdown automata, such instructions are independent from the ranks. To maintain the hierarchical character of the countdown, we choose a syntax that guarantees it by design: a *stack* of counters. Such alternative model, as well as effective language-preserving translations, are given in Section 4.7.

Countdown Complexity. We analyze complexity of definable languages. Given a language L definable in $\mu^{<\infty}$-ML, one can ask about the minimal nesting of the countdown operators required in a formula to define L. Another number related to L is the least stack height necessary for an automaton with stacked countdown to recognize L. We also introduce a syntactic parameter for countdown automata called countdown depth. Countdown depth leads to yet another measure of complexity: the least countdown depth necessary to recognize the language. Theorem 4.7.8 says that the three measures all coincide. The resulting parameter of L is called its *countdown complexity*. This stratifies languages into a *countdown hierarchy* of classes with greater and greater complexity. In Theorem 4.8.1 we show that under mild assumptions such hierarchy is strict.

Decidability Issues. Over finite models, the countdown operators do not introduce any new expressive power to μ-ML. This makes the (finite) model checking problem for $\mu^{<\infty}$-ML decidable but also less interesting. An arguably more interesting question here is that of *satisfiability*. We formulate Conjecture 4.9.2 according to which the satisfiability problem is decidable for the full logic $\mu^{<\infty}$-ML. Unfortunately, the lack of positional determinacy in countdown games prevents us from using proof techniques known from parity automata (where one can transform an alternating automaton into a nondeterministic one that guesses the positional strategy). Still, we use automata to solve the logic in a special case: Büchi automata and infinite words. These are countdown automata with only two ranks $r^{\forall} > r^{\exists}$. Theorem 4.9.5 says that satisfiability of Büchi automata over infinite words is decidable.

The full Conjecture 4.9.2 remains open. Nevertheless, the existence of an automata model equivalent to logic is encouraging. Apart from allowing us to solve some fragments of the logic, it implies that $\mu^{<\infty}$-ML does not share some of the troublesome properties of $\mathsf{MSO}+\mathsf{U}$ that result in undecidability. In particular, it can be used to show that all languages definable in μ-ML have *bounded topological complexity* (i.e. at most Σ^1_2, see [30] for an introduction to topological complexity in computer science). Since $\mathsf{MSO}+\mathsf{U}$ defines a Σ^1_n-complete language for every $n < \omega$ [18, Theorem 5.1], by [30, Theorem 7], it follows that some $\mathsf{MSO}+\mathsf{U}$-definable languages are not expressible in $\mu^{<\infty}$-ML (whether $\mu^{<\infty}$-ML-definability implies $\mathsf{MSO}+\mathsf{U}$-definability remains an open question). Since by [10, Theorem 1.3] every logic closed under boolean combinations, projections and defining the language from Example 4.2.5 contains $\mathsf{MSO}+\mathsf{U}$, this means that our calculus is *not closed under projections* and as a consequence *does not have an equivalent nondeterministic automata model.* This is an arguably good news, as in light of [5, Theorem 1.4], giving up closure under projections is the only way to go if one wants to design a decidable extension of MSO closed under boolean operations. Decidability of the weak variant $\mathsf{WMSO}+\mathsf{U}$ of $\mathsf{MSO}+\mathsf{U}$ over infinite words [3] and infinite (ranked) trees [7] shows that such extensions are possible. In fact, both results are obtained by establishing a correspondence with equivalent automata models, namely deterministic max-automata [3, Theorem 1] and nested limsup automata [7, Theorem 2]. Since the existence of accepting runs for such automata can be expressed in $\mu^{<\infty}$-ML, we get that $\mu^{<\infty}$-ML contains $\mathsf{WMSO}+\mathsf{U}$ on infinite words and trees. The opposite inclusion is false (due to topological reasons), at least for the trees. The relation between

$\mu^{<\infty}$-ML and the ω-B-, ω-S- and ω-BS-automata of [9] remains unclear, as these models do not admit determinization. Also, the relation between our logic and regular cost functions (see e.g. [12]) is less immediate than it could seem at first glance and requires further research.

History and Credits. The key mechanism of countdown games is implicit in [15], where the authors investigate a nonstandard semantics for the scalar fragment of the μ-calculus equivalent to replacing every μ and ν by our countdown operators μ^α and ν^α, respectively. However, the authors did not abstract from formulae in their definition of games, nor consider the full vectorial calculus that corresponds to automata. Countdown logic, automata and games in its mature form were introduced in [23] and the material presented in this chapter is an extended version of it. There are two notable exceptions: Section 4.7 and Subsection 4.9.1 present an entirely new material.

4.1 Games for Fixpoint Approximations

We have seen in Subsection 2.2.3 how the least fixpoint $\mathsf{LFP}.f$ and the greatest fixpoint $\mathsf{GFP}.f$ of a given monotone operation $f : \mathcal{P}(X) \to \mathcal{P}(X)$ can be characterized with games $\mathcal{G}_\mu(f)$ and $\mathcal{G}_\nu(f)$, respectively. According to the Knaster-Tarski Theorem 2.1.1, these fixpoints are computed as limits f_μ^∞ and f_ν^∞ of (transfinite) sequences $(f_\mu^\alpha)_{\alpha\in\mathsf{Ord}}$ and $(f_\nu^\alpha)_{\alpha\in\mathsf{Ord}}$ of their respective approximations. We start this chapter showing how the Definition 2.2.7 of fixpoint games can be modified to characterize these approximations of fixpoints. The idea is that $\mathcal{G}_\mu(f), x$ is a reachability game and so $\exists$ve is supposed to show that x belongs to f_μ^∞ in finitely many steps. In case of f_μ^α $\exists$ve's job should be even harder, as she is supposed to show that x is already included in the α-th approximation $f_\mu^\alpha \subseteq f_\mu^\infty$ of the fixpoint. To capture this intuition we enrich the game with a *counter* storing an ordinal value. At the beginning of each round $\exists$ve will decrement the counter and in case it reaches 0 she will loose. Symmetrically, to characterize f_ν^α we enrich $\mathcal{G}_\nu(f)$ with a counter decremented by $\forall$dam.

Definition 4.1.1. Fix a monotone $f : \mathcal{P}(X) \to \mathcal{P}(X)$. The games $\mathcal{G}_\mu^{<\infty}(f)$ and $\mathcal{G}_\nu^{<\infty}(f)$ are played in three-step rounds.

1. the *owner* of the counter, which is $\exists$ve in case of μ and $\forall$dam in case of ν, chooses a new counter value $\beta < \alpha$ smaller than the current value α (in particular, if $\alpha = 0$ the player is stuck and looses);

2. from the current position x $\exists$ve chooses $Y \in \mathcal{P}(X)$ such that $x \in f(Y)$;

3. $\forall$dam chooses $y \in Y$ and the next round starts from y with counter value β.

Formally:

$$\mathsf{Conf} = X \times \mathsf{Ord} \times \{\mathsf{cdn}, \mathsf{psn}\} \cup \mathcal{P}(X) \times \mathsf{Ord}$$

and:

$$\begin{aligned}\mathsf{Mov} =& \{((x, \alpha, \mathsf{cdn}), (x, \beta, \mathsf{psn})) \mid \beta < \alpha\} \cup \\ & \{((x, \alpha, \mathsf{psn}), (Y, \alpha)) \mid x \in f(Y)\} \cup \\ & \{((Y, \alpha), (x, \alpha, \mathsf{cdn})) \mid x \in Y\}.\end{aligned}$$

Configurations in $X \times \mathsf{Ord} \times \{\mathsf{psn}\}$ always belong to $\exists$ve, these in $\mathcal{P}(X) \times \mathsf{Ord}$ always belong to $\forall$dam and the ones from $X \times \mathsf{Ord} \times \{\mathsf{cdn}\}$ belong to $\exists$ve in case of μ and to $\forall$dam in case of ν. Since at each round the counter decreases, by well-foundedness of Ord there are no infinite plays and so the winning condition is insubstantial. We will write $\mathcal{G}^{\alpha}_{\mu}$ and $\mathcal{G}^{\alpha}_{\nu}$ to denote games with a default initial counter value $\alpha \in \mathsf{Ord}$.

The game is designed so that it captures the approximations of fixpoints in analogy to Theorem 2.2.8.

Theorem 4.1.2. *For every $x \in X$ and $\alpha \in \mathsf{Ord}$:*

1. *$\exists$ve wins $\mathcal{G}^{<\infty}_{\mu}(f)$ from $(x, \alpha, \mathsf{cdn})$ iff $x \in f^{\alpha}_{\mu}$.*
2. *$\exists$ve wins $\mathcal{G}^{<\infty}_{\nu}(f)$ from $(x, \alpha, \mathsf{cdn})$ iff $x \in f^{\alpha}_{\nu}$.*

Proof. The proof proceeds by immediate transfinite induction on α. Assume that the claim is true for all $\beta < \alpha$. For both $\eta = \mu$ and $\eta = \nu$, it follows from the respective items of Proposition 2.2.9 that $\exists$ve has a strategy for the round starting at $(x, \alpha, \mathsf{cdn})$ only leading to winning configurations iff $x \in f^{\alpha}_{\eta}$. $\square$

We have illustrated the game characterization of fixpoints from Theorem 2.2.8 with Example 2.2.10. Using the theorem we have derived adequacy (2.2) of the bisimulation game $\mathcal{G}_{\leftrightarroweq}(\mathcal{M}, \mathcal{M}')$ from the fact that bisimilarity $\leftrightarroweq$ is the greatest fixpoint of the operation $\mathsf{BIS} : \mathcal{P}(M \times M') \to \mathcal{P}(M \times M')$. Let us now demonstrate how an analogous characterization for depth-k bisimilarity $\leftrightarroweq^k$ follows from Theorem 4.1.2.

Example 4.1.3. We show that adequacy (2.3) of the game $\mathcal{G}^{<\omega}_{\leftrightarroweq}(\mathcal{M}, \mathcal{M}')$ for depth-k bisimilarity $\leftrightarroweq^k$ follows from Theorem 4.1.2. To that end, observe that for all $k < \omega$ the relation $\leftrightarroweq^{k+1}$ is precisely $\mathsf{BIS}(\leftrightarroweq^k)$. Since $\leftrightarroweq^0$ is the full relation it follows that:

$$\leftrightarroweq^k = \mathsf{BIS}^k_{\nu}$$

for all $k < \omega$. Hence, the theorem implies that $\exists$ve wins $\mathcal{G}^{<\infty}_{\nu}(\mathsf{BIS})$ from $(\mathsf{m}, \mathsf{m}', k, \mathsf{cdn})$ iff $\mathsf{m} \leftrightarroweq^k \mathsf{m}'$. Therefore, to prove (2.3) it suffices to show:

$$\mathcal{G}^{<\infty}_{\nu}(\mathsf{BIS}), ((\mathsf{m}, \mathsf{m}'), k, \mathsf{cdn}) \;\bowtie\; \mathcal{G}^{<\omega}_{\leftrightarroweq}(\mathcal{M}, \mathcal{M}'), (\mathsf{m}, \mathsf{m}', k, \mathsf{count}) \tag{4.1}$$

for all $\mathsf{m} \in M, \mathsf{m}' \in M'$ and $k < \omega$.

Proof. To prove the equivalence, first we relax the definition of $\mathcal{G}_{\leftrightarrows}^{<\omega}$ so that in the countdown step count instead of a deterministic decrement of the counter (leading from k to $k-1$ or to $\exists$ve's victory in case of value 0) we ask $\forall$dam to choose its new smaller value. Since bigger values are always better for $\forall$dam, such modification does not change the winner of the game. Moreover, after such a relaxation there is nothing special about finite numbers and we may generalize $\mathcal{G}_{\leftrightarrows}^{<\omega}$ by allowing the counter to take arbitrary ordinal values. Denote such generalized game by $\mathcal{G}_{\leftrightarrows}^{<\infty}$. Because the counter value never increases, if $\mathcal{G}_{\leftrightarrows}^{<\infty}$ starts with finite one it is equivalent to the original game $\mathcal{G}_{\leftrightarrows}^{<\omega}$ initialized in the same configuration. Hence, for (4.1) it suffices if we prove:

$$\mathcal{G}, ((\mathsf{m}, \mathsf{m}'), \alpha, \mathsf{cdn}) \;\bowtie\; \mathcal{G}', (\mathsf{m}, \mathsf{m}', \alpha, \mathsf{count}) \tag{4.2}$$

for all $\mathsf{m} \in M, \mathsf{m}' \in M'$ and $\alpha \in \mathsf{Ord}$ where:

$$\mathcal{G} = \mathcal{G}_{\nu}^{<\infty}(\mathsf{BIS}) \qquad \text{and} \qquad \mathcal{G}' = \mathcal{G}_{\leftrightarrows}^{<\infty}(\mathcal{M}, \mathcal{M}').$$

Towards the use of the Decomposition Lemma 2.2.5 we view both $\mathcal{G}$ and $\mathcal{G}'$ as parity games in a similar fashion as in Example 2.2.10: we assign the same rank belonging to $\forall$dam to all the configurations. Note that in this case the assignment of ranks is purely technical and does not matter for the games as neither game allows for infinite plays. It reflects, however, that $\forall$dam wants to end the game within as few rounds as possible. Consider configurations $\mathcal{S}_+ \subseteq \mathsf{Conf}$ and $\mathcal{S}'_+ \subseteq \mathsf{Conf}'$ from which each round starts i.e.:

$$\mathcal{S}_+ = (M \times M') \times \mathsf{Ord} \times \{\mathsf{cdn}\} \qquad \text{and} \qquad \mathcal{S}'_+ = M \times M' \times \mathsf{Ord} \times \{\mathsf{count}\}.$$

With such rank functions we apply the Decomposition Lemma 2.2.5 (substituting $S = \emptyset$) and deduce (4.2) from:

$$\mathcal{G}, ((\mathsf{m}, \mathsf{m}'), \alpha, \mathsf{cdn}) \;\bowtie_{S_+}\; \mathcal{G}', (\mathsf{m}, \mathsf{m}', \alpha, \mathsf{count})$$

where $S_+ \subseteq \mathcal{S}_+ \times \mathcal{S}'_+$ is the relation that links configurations $((\mathsf{m}, \mathsf{m}'), \alpha, \mathsf{cdn})$ and $(\mathsf{m}, \mathsf{m}', \alpha, \mathsf{count})$ that have the same points and counter value:

$$S_+ = \{(((\mathsf{n}, \mathsf{n}'), \alpha, \mathsf{cdn}), (\mathsf{n}, \mathsf{n}', \alpha, \mathsf{count})) \mid \mathsf{n} \in M, \mathsf{n}' \in M', \alpha \in \mathsf{Ord}\}.$$

Since both games start with $\forall$dam choosing a new counter value $\beta < \alpha$ and the stage changes to psn and base, respectively, it suffices if we prove:

$$\mathcal{G}, ((\mathsf{m}, \mathsf{m}'), \beta, \mathsf{psn}) \;\bowtie_{\mathcal{S}_+}\; \mathcal{G}', (\mathsf{m}, \mathsf{m}', \beta, \mathsf{base}).$$

This follows from composing:

$$\begin{aligned}\mathcal{G}, ((\mathsf{m}, \mathsf{m}'), \beta, \mathsf{psn}) &\bowtie_{S_\beta} \mathcal{G}_{\nu}(\mathsf{BIS}), (\mathsf{m}, \mathsf{m}') \\ &\bowtie_{R} \mathcal{G}_{\leftrightarrows}, (\mathsf{m}, \mathsf{m}', \mathsf{base}) \\ &\bowtie_{S'_\beta} \mathcal{G}', (\mathsf{m}, \mathsf{m}', \beta, \mathsf{base})\end{aligned}$$

where:

$$\begin{aligned}
S_\beta &= \{(((\mathsf{n},\mathsf{n}'),\beta,\mathsf{cdn}),(\mathsf{n},\mathsf{n}')) \mid \mathsf{n}\in M, \mathsf{n}'\in M'\},\\
R &= \{((\mathsf{n},\mathsf{n}'),(\mathsf{n},\mathsf{n}',\mathsf{base})) \mid \mathsf{n}\in M, \mathsf{n}'\in M'\},\\
S'_\beta &= \{((\mathsf{n},\mathsf{n}',\mathsf{base}),(\mathsf{n},\mathsf{n}',\beta,\mathsf{count}) \mid \mathsf{n}\in M, \mathsf{n}'\in M'\}.
\end{aligned}$$

The first equivalence follows from the observation that the counter will not be modified before $\mathcal{G}$ stops; similarly with the third equivalence and $\mathcal{G}'$. The middle one is (2.7). The composition:

$$S_\beta \circ R \circ S'_\beta = \{(\gamma,\gamma') \in S_+ \mid \text{the counters have value } \beta\}$$

is precisely S_+ restricted to configurations with counter value β. Since all the configurations accessible in $\mathcal{G}$ and $\mathcal{G}'$ before stopping have unchanged counter value equal to β, this completes the proof. □

Note that although we mostly focus on finite-depth bisimilarity, depth-α bisimilarity for arbitrary (finite or infinite) ordinal α could be defined as:

$$\leftrightarroweq^\alpha = \mathsf{BIS}^\alpha_\nu.$$

Then, the above proof shows that such a relation is characterized with $\mathcal{G}^{<\infty}_{\leftrightarroweq}(\mathcal{M},\mathcal{M}')$:

$$\exists\text{ve wins } \mathcal{G}^{<\omega}_{\leftrightarroweq}(\mathcal{M},\mathcal{M}'),(\mathsf{m},\mathsf{m}',\alpha,\mathsf{count}) \iff \mathcal{M},\mathsf{m} \leftrightarroweq^\alpha \mathcal{M}',\mathsf{m}'.$$

for all $\alpha \in \mathsf{Ord}$.

Interestingly enough, history comes full circle here. Initially, $\leftrightarroweq^\omega$ was considered as a candidate for the notion capturing behavioral equivalence [28]. When $\leftrightarroweq$ was introduced it replaced the more complicated $\leftrightarroweq^k$. Now, we return from the smooth theory of fixpoints and related games back to their approximations in the hope of broadening our understanding.

4.2 Countdown Logic

We now introduce the *countdown μ-calculus* $\mu^{<\infty}$-ML. We begin with the scalar version.

4.2.1 The Scalar Fragment

As before, fix an infinite set Var of variables and a set Act of actions. The syntax of *(scalar) countdown μ-calculus* is defined as follows:

$$\varphi ::= \top \mid \bot \mid \tau \mid \neg\tau \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \langle \mathsf{a}\rangle\varphi \mid [\mathsf{a}]\varphi \mid x \mid \mu^\alpha x.\varphi \mid \nu^\alpha x.\varphi \quad (4.3)$$

with $\tau \in \mathsf{Prop}$, $\mathsf{a} \in \mathsf{Act}$, $x \in \mathsf{Var}$ and $\alpha \in \mathsf{Ord}_\infty$; the presence of ordinal numbers α is the only syntactic difference with the standard μ-ML.

The semantics of countdown formulae is defined inductively the same way as for the standard μ-ML with the only difference that μ^α and ν^α are interpreted as F^α_μ and F^α_ν instead of F^∞_μ and F^∞_ν:

$$
\begin{aligned}
[\![\top]\!]^{\mathsf{val}} = M &\quad \text{and} \quad [\![\bot]\!]^{\mathsf{val}} = \emptyset \\
[\![\tau]\!]^{\mathsf{val}} = \tau^{\mathcal{M}} &\quad \text{and} \quad [\![\neg\tau]\!]^{\mathsf{val}} = M - \tau^{\mathcal{M}} \\
[\![\varphi_1 \vee \varphi_2]\!]^{\mathsf{val}} = [\![\varphi_1]\!]^{\mathsf{val}} \cup [\![\varphi_2]\!]^{\mathsf{val}} &\quad \text{and} \quad [\![\varphi_1 \wedge \varphi_2]\!]^{\mathsf{val}} = [\![\varphi_1]\!]^{\mathsf{val}} \cap [\![\varphi_2]\!]^{\mathsf{val}}; \\
[\![\langle \mathsf{a}\rangle\varphi]\!]^{\mathsf{val}} = \{\mathsf{m} \in M \mid \exists_{\mathsf{m} \xrightarrow{\mathsf{a}} \mathsf{n}}\ \mathsf{n} \in [\![\varphi]\!]^{\mathsf{val}}\} &\quad \text{and} \quad [\![[\mathsf{a}]\varphi]\!]^{\mathsf{val}} = \{\mathsf{m} \in M \mid \forall_{\mathsf{m} \xrightarrow{\mathsf{a}} \mathsf{n}}\ \mathsf{n} \in [\![\varphi]\!]^{\mathsf{val}}\} \\
[\![x]\!]^{\mathsf{val}} &= \mathsf{val}(x) \\
[\![\mu^\alpha x.\varphi]\!]^{\mathsf{val}} = F^\alpha_\mu &\quad \text{and} \quad [\![\nu^\alpha x.\varphi]\!]^{\mathsf{val}} = F^\alpha_\nu
\end{aligned}
$$

where in the last clause $F(H) = [\![\varphi]\!]^{\mathsf{val}[x \mapsto H]}$.

This contains the classical μ-calculus, with μ^∞ and ν^∞ being just the ordinary μ and ν fixpoint operators. However, $\mu^{<\infty}$-ML is capable of capturing *boundedness* and *unboundedness* properties which are not expressible in the classical setting:

Example 4.2.1. For $|\mathsf{Act}| = 1$, consider the formula $\mu^\alpha x.\Box x$ and compare it to the fixpoint formula $\mu x.\Box x$ from Example 2.4.1. The semantics of both formulae in a model $\mathcal{M}$ are obtained by iterating the same map:

$$H \overset{F}{\mapsto} \{\mathsf{m} \in M \mid \forall_{\mathsf{m} \to \mathsf{n}}.\ \mathsf{n} \in H\}.$$

The only difference is that in the case of μ the map is iterated until reaching a fixpoint F^∞_μ, whereas with μ^ω the process stops after ω steps with F^ω_μ. For $\alpha < \omega$ the set F^α_μ consists of the points from which there is no path longer than α. Hence, $[\![\mu^\omega x.\Box x]\!] = F^\omega_\mu$ is the set of all points for which there exists a finite bound on the lengths of paths starting there. Example 2.1.2 illustrates that.

4.2.2 The Vectorial Calculus

The (full) *countdown μ-calculus* is defined as for its scalar fragment, except that fixpoint operators act on tuples (vectors) of formulae rather than on single formulae.

Definition 4.2.2. The syntax of *countdown μ-calculus* is given as follows:

$$\varphi ::= \top \mid \bot \mid \tau \mid \neg\tau \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \langle \mathsf{a}\rangle\varphi \mid [\mathsf{a}]\varphi \mid x \mid \mu^\alpha_i \overline{x}.\overline{\varphi} \mid \nu^\alpha_i \overline{x}.\overline{\varphi}$$

where $\tau \in \mathsf{Prop}$, $\mathsf{a} \in \mathsf{Act}$, $x \in \mathsf{Var}$ and $\alpha \in \mathsf{Ord}_\infty$ as in the scalar fragment and additionally $1 \leq i \leq d < \omega$ with $\overline{x} = \langle x_1, ..., x_d\rangle \in \mathsf{Var}^d$ and $\overline{\varphi} = \langle \varphi_1, ..., \varphi_d\rangle$ being d-tuples of variables and formulae, respectively.

Definition 4.2.3. The meaning $[\![\varphi]\!]^{\mathsf{val}} \subseteq M$ of a formula φ in a model $\mathcal{M}$ under valuation val is defined by induction the same way as for the scalar formulae except for the operators μ_i^α and ν_i^α, in which case:

$$[\![\mu_i^\alpha \overline{x}.\overline{\varphi}]\!]^{\mathsf{val}} = \pi_i(F_\mu^\alpha) \quad \text{and} \quad [\![\nu_i^\alpha \overline{x}.\overline{\varphi}]\!]^{\mathsf{val}} = \pi_i(F_\nu^\alpha)$$

where the monotone map $F : (\mathcal{P}(M))^d \to (\mathcal{P}(M))^d$ is given as:

$$F \begin{pmatrix} H_1 \\ \vdots \\ H_d \end{pmatrix} = \begin{pmatrix} [\![\varphi_1]\!]^{\mathsf{val}'} \\ \vdots \\ [\![\varphi_d]\!]^{\mathsf{val}'} \end{pmatrix}$$

for $\mathsf{val}' = \mathsf{val}[x_1 \mapsto H_1, ..., x_d \mapsto H_d]$ and $\pi_i : (\mathcal{P}(M))^d \to \mathcal{P}(M)$ is the i-th projection.

Note that operators μ^∞ and ν^∞ are equivalent to μ and ν from the classical μ-calculus. Furthermore, for every ordinal α, the formula $\mu_i^{\alpha+1}\overline{x}.\overline{\psi}$ is equivalent to:

$$\psi_i[x_1 \mapsto \mu_1^\alpha \overline{x}.\overline{\psi}, \ldots, x_d \mapsto \mu_d^\alpha \overline{x}.\overline{\psi}]$$

and similarly for $\nu^{\alpha+1}$. As a result, without loss of generality we may assume that in countdown operators μ^α and ν^α only limit ordinals α are used.

The countdown μ-calculus is semantically closed under negation in the same way as the classical calculus: for every formula φ one can construct $\widetilde{\varphi}$ dual to φ, meaning that:

$$[\![\widetilde{\varphi}]\!]^{\mathsf{val}} = M - [\![\varphi]\!]^{\widetilde{\mathsf{val}}}$$

for every model $\mathcal{M}$ and valuations val and $\widetilde{\mathsf{val}}$ such that $\widetilde{\mathsf{val}}(x) = M - \mathsf{val}(x)$ for all $x \in \mathsf{Var}$. The standard inductive definition is then extended with straightforward:

$$\widetilde{\mu_i^\alpha \overline{x}.\overline{\varphi}} = \nu_i^\alpha \overline{x}.\overline{\widetilde{\varphi}} \quad \text{and} \quad \widetilde{\nu_i^\alpha \overline{x}.\overline{\varphi}} = \mu_i^\alpha \overline{x}.\overline{\widetilde{\varphi}}.$$

Example 4.2.4. Consider the formula $\mu^\omega x.\Box x$ from Example 4.2.1 defining boundedness. It follows that its dual $\widetilde{\mu^\omega x.\Box x} = \nu^\omega x.\Diamond x$ is true at a point if there are arbitrarily long paths starting there. Since the formula $\mu x.\Box x$ from Example 2.4.1 describes well-foundedness, the conjunction $\nu^\omega x.\Diamond x \wedge \mu x.\Box x$ is true at a point iff there are *arbitrarily long* paths starting there but no *infinite* one. Although such a formula is satisfiable, for example in The Hedgehog from Example 3.1.2, it follows by the König's Lemma that every its model must be infinite. This demonstrates that, unlike μ-ML, $\mu^{<\infty}$-ML does not have the finite model property.

In Section 4.6 we will compare the expressive power of the vectorial and scalar countdown μ-calculus in detail. For now, let us show that the Bekić principle (2.9) fails for countdown operators:

Example 4.2.5. An infinite word $W \in \Gamma^\omega$ over the alphabet $\Gamma = \{\mathsf{a}, \mathsf{b}\}$ can be seen as a model for $\mathsf{Act} = \Gamma$ with ω as the set of points and with transition relations defined by:

$$n \xrightarrow{\mathsf{a}} m \iff m = n + 1 \text{ and } W_n = \mathsf{a}.$$

For every regular language $K \subseteq \Gamma^*$ and $x \in \mathsf{Var}$, it is straightforward to define a fixpoint formula (in the classical μ-calculus, so without countdown operators) $\langle K \rangle x$ that holds in a point n, for a valuation val, if and only if there exists a word $w \in K$ and a path in W labelled with w that starts in n and ends in a point that belongs to $\mathsf{val}(x)$. Then, the formula:

$$\varphi = \nu_1^\omega(x_1, x_2).(\langle \Gamma^* \rangle x_2, \langle \mathsf{a} \rangle x_2)$$

is true in a word W iff it contains arbitrarily long blocks of consecutive a's. To see this, observe that at the i-th step of approximation: (i) the second component (x_2) contains a point n iff the next i transitions are all labelled with a, and (ii) the first component (x_1) contains a point n iff the second component contains at least one point after n.

However, the following scalar formula constructed by analogy to the Bekić principle:

$$\psi = \nu^\omega x_1.\langle \Gamma^* \rangle (\nu^\omega x_2.\langle \mathsf{a} \rangle x_2)$$

is equivalent to $\langle \Gamma^* \rangle (\nu^\omega x_2.\langle \mathsf{a} \rangle x_2)$, and the formula under $\langle \Gamma^* \rangle$ holds in a point iff all the future transitions from that point are labelled with a. Thus, ψ holds (in any point) iff the word W is of the form $\Gamma^* \mathsf{a}^\omega$, and so ψ is not equivalent to φ.

Let us emphasize that the above counterexample to a principle analogous to the Bekić rule relies on the fact that the index ω in the operator ν^ω is an ordinal, not ∞. The operators μ^α and ν^α with $\alpha = \infty$ are ordinary fixpoint operators and so the principle (2.9) allows us to rewrite every countdown formula to a form where the only non-scalar operators are μ^α and ν^α with $\alpha \neq \infty$.

A number of good properties of μ-ML comes from the tight connection between logic, games and automata. As it turns out, such a threefold correspondence can be lifted to the countdown setting if we appropriately modify the definition of a parity game and then consider *countdown automata* arising from such *countdown games*.

4.3 Countdown Games

To match with the classical μ-ML we needed a parity winning condition: the logic involving nesting of fixpoint operators corresponds to the parity condition thought of as a nested safety/reachability condition. The countdown calculus $\mu^{<\infty}$-ML allows for nesting of countdown operators μ^α and

ν^α. Since the semantics of each such η^α is captured by respective game $\mathcal{G}_\eta^{<\infty}$ featuring a single counter, this naturally leads to the notion of *countdown games* extending simple parity games with multiple counters. The counters are organized in a hierarchical manner reflecting the hierarchical character of nesting.

Intuitively, the counters say how many more times various ranks can be visited, in similar manner to the signatures introduced by Walukiewicz [38, Section 3]. A player responsible for decrementing a counter may lose the game if the value of that counter is zero, just as a player responsible for finding the next position in a game may lose if there is no position to go to.

Definition 4.3.1. Syntactically, a countdown game is given as a tuple $(V, E, \mathsf{rank}, \mathsf{ctr}_I)$ such that (V, E, rank) is a simple parity game and $\mathsf{ctr}_I : \mathcal{D} \to \mathsf{Ord}$ is a map from a subset $\mathcal{D} \subseteq \mathcal{R}$ of ranks to ordinals. We call ctr_I the *initial counter assignment* and the set $\mathcal{D}$ *nonstandard ranks* of the game. The idea is that nonstandard ranks have associated counters and at positions with such ranks countdown will occur.

Explicitly, a countdown game assumes:

- a set $V = V_\exists \sqcup V_\forall$ of positions,
- an edge relation $E \subseteq V \times V$, and
- a rank function $\mathsf{rank} : V \to \mathcal{R}$ for a fixed finite linear order $\mathcal{R} = \mathcal{R}_\exists \sqcup \mathcal{R}_\forall$.

In addition, we fix:

- a subset $\mathcal{D} \subseteq \mathcal{R}$ of *nonstandard* ranks and
- an *initial counter assignment* $\mathsf{ctr}_I : \mathcal{D} \to \mathsf{Ord}$.

Unlike with simple parity games, configurations of a countdown game are not identified with its mere positions. Each configuration consists of a position $v \in V$, a counter assignment $\mathsf{ctr} : \mathcal{D} \to \mathsf{Ord}$, and a bit of information from the set $\{\mathsf{cdn}, \mathsf{psn}\}$: [1]

$$\mathsf{Conf} = V \times \mathsf{Ord}^{\mathcal{D}} \times \{\mathsf{cdn}, \mathsf{psn}\}.$$

The last component encodes one of the two possible kinds of configurations, called *countdown* and *positional* configurations. Positional configurations $(v, \mathsf{ctr}, \mathsf{psn})$ are owned by the owner of v, and countdown ones $(v, \mathsf{ctr}, \mathsf{cdn})$ by the owner of its rank $\mathsf{rank}(v)$. The possible moves $\mathsf{Mov} \subseteq \mathsf{Conf} \times \mathsf{Conf}$ are given as follows:

[1] A reader concerned about set-theoretic issues can modify the above definition by replacing arbitrary counter assignments $\mathsf{Ord}^{\mathcal{D}}$ with only those that have all values bounded by the initial ones $\{\mathsf{ctr} \in \mathsf{Ord}^{\mathcal{D}} \mid \mathsf{ctr} \leq \mathsf{ctr}_I\}$ so that Conf is a proper set. Since only the latter assignments are reachable, the two definitions are equivalent.

- From a countdown configuration $(v, \mathsf{ctr}, \mathsf{cdn})$, the owner of $r = \mathsf{rank}(v)$ chooses a counter assignment ctr' such that:
  - $\mathsf{ctr}'(r') = \mathsf{ctr}_I(r')$ for $r' < r$,
  - $\mathsf{ctr}'(r) < \mathsf{ctr}(r)$ (if r is nonstandard),
  - $\mathsf{ctr}'(r') = \mathsf{ctr}(r')$ for $r' > r$,

  and the game proceeds from the positional configuration $(v, \mathsf{ctr}', \mathsf{psn})$.

  In words: counters for ranks lower than r are reset, the counter for r (if any) is decremented, and counters for higher ranks are left unchanged. Note that if r is standard then there is no real choice here: ctr' is determined by ctr. And if r is nonstandard then the move amounts to choosing an ordinal $\alpha < \mathsf{ctr}(r)$.

- From a positional configuration $(v, \mathsf{ctr}, \mathsf{psn})$, the owner of v chooses an edge $(v, w) \in E$ and the game proceeds from the countdown configuration $(w, \mathsf{ctr}, \mathsf{cdn})$.

In any configuration, if the player responsible for making the next move is stuck, (s)he looses immediately. Otherwise, in an infinite play, the winner is determined by the parity condition: the owner of the greatest rank appearing infinitely often looses.

The default initial counter assignment is ctr_I and the default initial mode is the countdown one, meaning that $\mathcal{G}, v$ stands for $\mathcal{G}, (v, \mathsf{ctr}_I, \mathsf{cdn})$. We denote $\mathcal{D}_\exists = \mathcal{D} \cap \mathcal{R}_\exists$ and $\mathcal{D}_\forall = \mathcal{D} \cap \mathcal{R}_\forall$. In the context of arbitrary parity games we denote the lowest irrelevant rank with 0. In the particular case of countdown games, we will always assume that this rank is standard, so that the counter update corresponding to 0 is trivial (i.e. no counter changes).

Both countdown games and games for fixpoint approximations involve the positional and countdown modes, marked with psn and cdn, respectively. This is not a coincidence: as we mentioned, countdown games generalize the games for fixpoint approximations the same way as simple parity games generalize games for fixpoints. Let us inspect this connection in a bit more detail.

Example 4.3.2. Fix a monotone operation $f : \mathcal{P}(X) \to \mathcal{P}(X)$. We show how the game $\mathcal{G}^\alpha_\mu(f)$ for the α-th approximation of the least fixpoint (as given in Definition 4.1.1) arises from $\mathcal{G}_\mu(f)$ (the case with $\mathcal{G}^\alpha_\nu(f)$ and the α-th approximation of the greatest fixpoint is analogous). Recall that according to Definition 2.2.7 the game $\mathcal{G}_\mu(f)$ is played in rounds consisting of two alternating steps:

1. from $x \in X$, $\exists$ve comes up with $Y \subseteq X$ such that $x \in f(Y)$;
2. $\forall$dam chooses $y \in Y$ and the next round starts from there.

We want to view $\mathcal{G}_\mu(f)$ as a simple parity game $\mathcal{G} = (V, E, \mathsf{rank})$. The positions V are just the configurations of $\mathcal{G}_\mu(f)$ and edges $E \subseteq V \times V$ are its legal moves. Since $\mathcal{G}_\mu(f)$ is a reachability game, it could seem natural to assign one rank r belonging to ∃ve to all the positions. However, we only assign r to the positions in X and leave the remaining positions from $\mathcal{P}(X)$ with an irrelevant standard rank $0 < r$. This does not change the winner of any play, but reflects the fact that we want to count the number of rounds, not moves. With $\mathcal{G}$ defined this way, we obtain $\mathcal{G}^\alpha_\mu(f)$ by turing r into a nonstandard rank with initial counter value α. The resulting game $\mathcal{G}' = (V, E, \mathsf{rank}, \mathsf{ctr}_I)$ is played in rounds consisting of three alternating steps:

1. from configuration consisting of a position $x \in X$ and a counter value β, ∃ve decrements the counter by picking some $\beta' < \beta$;
2. next, ∃ve comes up with $Y \subseteq X$ such that $x \in f(Y)$;
3. ∀dam chooses $y \in Y$ and the next round starts from (y, β').

In the above description we skip the trivial counter updates corresponding to rank 0. It follows that the games $\mathcal{G}'$ and $\mathcal{G}^\alpha_\mu(f)$ are isomorphic (up to skipping the trivial counter updates).

The case with ν in place of μ is the same except that the rank r belongs to ∀dam and thus if we turn it into a nonstandard one it is him who controls the counter. In Example 4.1.3 we showed how the game $\mathcal{G}^{<\infty}_{\Leftrightarrow}(\mathcal{M}, \mathcal{M}')$ characterizing depth-α bisimilarity $\Leftrightarrow^\alpha \subseteq M \times M'$ is the same as $\mathcal{G}^\alpha_\nu(\mathsf{BIS})$ for the operation $\mathsf{BIS} : \mathcal{P}(M \times M') \to \mathcal{P}(M \times M')$. In light of the above discussion, this allows us to see $\mathcal{G}^{<\infty}_{\Leftrightarrow}(\mathcal{M}, \mathcal{M}')$ as a countdown game. However, an explicit description similar to the one above can be easily given. Take $\mathcal{G}_{\Leftrightarrow}(\mathcal{M}, \mathcal{M}')$ and view it as a simple parity game with the round-beginning positions $M \times M' \times \{\mathsf{base}\}$ having the most important rank r belonging to ∀dam and all the other positions having an irrelevant lower rank 0. If we now turn that r into a nonstandard rank, we obtain $\mathcal{G}^\alpha_\nu(\mathsf{BIS})$ (again, up to skipping the trivial counter updates).

Every play of the game alternates between positional and countdown configurations, and in each move only one component of the configuration is modified. Therefore, although a play is formally a sequence of configurations, it can be more succinctly represented as an alternating sequence of positions and counter assignments:

$$\pi = \mathsf{ctr}_1 v_1 \mathsf{ctr}_2 v_2 \mathsf{ctr}_3 v_3 \cdots \tag{4.4}$$

Note that the only way the counters may interfere with a play is when a counter has value 0 and so its owner cannot decrement it. It is therefore beneficial for a player to have greater ordinals at his/her counters. More precisely, given a countdown game, define a partial order $\preccurlyeq_{\mathsf{better}}$ on its configurations: $(v, \mathsf{ctr}, \mathsf{psn}) \preccurlyeq_{\mathsf{better}} (v, \mathsf{ctr}', \mathsf{psn})$ and $(v, \mathsf{ctr}, \mathsf{cdn}) \preccurlyeq_{\mathsf{better}} (v, \mathsf{ctr}', \mathsf{cdn})$ if and only if $\mathsf{ctr}(r) \leq \mathsf{ctr}'(r)$ for all $r \in \mathcal{D}_\exists$ and $\mathsf{ctr}(r) \geq \mathsf{ctr}'(r)$ for all $r \in \mathcal{D}_\forall$. It easily follows from the definition that if $\gamma \preccurlyeq_{\mathsf{better}} \gamma'$ and $\exists$ve has a move from γ to a configuration δ, then she has a move from γ' to a δ' such that $\delta \preccurlyeq_{\mathsf{better}} \delta'$. Symmetrically, if $\forall$dam has a move from γ' to δ' then he has a move from γ to some δ with $\delta \preccurlyeq_{\mathsf{better}} \delta'$. That is: $\preccurlyeq_{\mathsf{better}}$ is a gamemulation order. As a result, if $\exists$ve has a winning strategy from γ and $\gamma \preccurlyeq_{\mathsf{better}} \gamma'$ then she has a winning strategy from γ'.

Another easy observation is that if the countdown starts from limit ordinals, one may always choose values greater by a finite k than the ones given by some fixed strategy. More specifically, for a number $k < \omega$ we define a relation above_{+k} of being k-*above* between counter assignments. For every ctr and ctr', $\mathsf{above}_{+k}(\mathsf{ctr}, \mathsf{ctr}')$ iff:

- $\mathsf{ctr}'(r) = \min(\mathsf{ctr}(r) + k, \mathsf{ctr}_I(r))$ for $r \in \mathcal{D}_\exists$, and
- $\mathsf{ctr}'(r) = \mathsf{ctr}(r)$ for $r \in \mathcal{D}_\forall$.

The relation above_{+k} extends to configurations in a natural way: $\mathsf{above}_{+k}(\gamma, \gamma')$ iff the only difference between γ and γ' is in their respective counter assignments ctr and ctr' with $\mathsf{above}_{+k}(\mathsf{ctr}, \mathsf{ctr}')$. It follows from the definition that if all initial counter values are limit ordinals then such above_{+k} is a gamemulation.

As all parity games, countdown games are configurationally determined. The following example demonstrates that, unlike *simple* parity games, countdown games are not *positionally* determined, in the sense that the players may need to look at the counter values in order to choose a winning move.

Example 4.3.3. Consider the Hedgehog from Example 3.1.2:

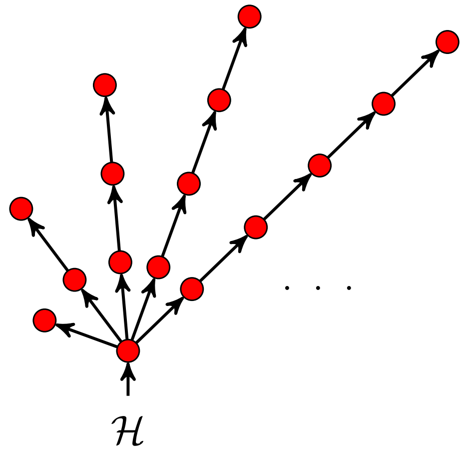

and view it as the arena (V, E) of a countdown game with all the positions $V = V_\exists$ belonging to $\exists$ve. Assign the same rank r belonging to $\forall$dam to all the positions and make r nonstandard with initial counter value ω. It follows that the game alternates between two stages: $\forall$dam decrementing the counter and $\exists$ve picking a move in the arena. Since at the beginning $\forall$dam needs to pick $k < \omega$, $\exists$ve can always win by choosing a path longer than k so that the counter will hit 0 before the game reaches a dead-end position. However, there is no uniform choice of a one path that would allow her to win against all k.

Even though not positional, countdown games possess a much weaker (yet still useful) property: players can win with strategies that do not depend on the counters in finite stages of the game. Consider a countdown game $(V, E, \mathsf{rank}, \mathsf{ctr}_I)$. For a countdown play:

$$\pi = \mathsf{ctr}_1 v_1 ... v_{n-1} \mathsf{ctr}_n v_n, \quad \text{or} \quad \pi = \mathsf{ctr}_1 v_1 ... \mathsf{ctr}_n v_n \mathsf{ctr}_{n+1}$$

denote by $\mathsf{pos}(\pi)$ the sequence of consecutive positions $v_1...v_n$. Recall that a phase of a game is a set of plays convex with respect to the prefix order. Given a phase $\mathcal{B}$ and a strategy σ for player P, we say that a partial function $f : V^* \to V$ *guides* σ in $\mathcal{B}$ if for every σ-plays $\pi, \pi v \in \mathcal{B}$ such that $v \in V$ is a position chosen by P, the value $f(\mathsf{pos}(\pi))$ is defined and equals v. We say that σ is *counter-independent in* $\mathcal{B}$ or $\mathcal{B}$*-counter-independent* iff it is guided in $\mathcal{B}$ by some partial function called the $\mathcal{B}$*-guide* of σ and denoted $\sigma^{\mathcal{B}}$. We call $\mathcal{B}$ *proper* if membership in $\mathcal{B}$ does not depend on the counter values, meaning that for plays π, π' of the same length, $\mathsf{pos}(\pi) = \mathsf{pos}(\pi')$ implies $\pi \in \mathcal{B} \iff \pi' \in \mathcal{B}$.

Proposition 4.3.4. *Take a countdown game $\mathcal{G} = (V, E, \mathsf{rank}, \mathsf{ctr}_I)$ and a proper phase $\mathcal{B}$ of $\mathcal{G}$. Assume that the set $\mathsf{pos}[\mathcal{B}] = \{\mathsf{pos}(\pi) \mid \pi \in \mathcal{B}\}$ is finite. If $\exists$ve wins from configuration γ_I, then she wins with a strategy that is counter-independent in $\mathcal{B}$.*

Proof. The assumption on $\mathsf{pos}[\mathcal{B}]$ implies that there exists a finite bound l_{max} on the length of plays in $\mathcal{B}$. Consider a winning strategy σ for $\exists$ve. We show by induction on $0 \le l \le l_{\mathsf{max}}$ that:

> For every σ-play π of length $|\pi| = l_{\mathsf{max}} - l$, there exists a winning strategy σ_π for $\mathcal{G}, \gamma_I$ that is counter-independent in the subphase $\mathcal{B}_\pi$ of $\mathcal{B}$ and equal to σ on plays without a prefix from $\mathcal{B}_\pi$.

Once we prove the claim for $l = l_{\mathsf{max}}$, we obtain a strategy σ_ε counter-independent in $\mathcal{B}_\varepsilon = \mathcal{B}$, as desired.

The base case is $l = 0$ where there is nothing to prove, as $|\pi| = l_{\mathsf{max}}$ implies that either $\mathcal{B}_\pi = \{\pi\}$ if $\pi \in \mathcal{B}$ or $\mathcal{B}_\pi = \emptyset$ otherwise. In both cases σ is trivially guided in $\mathcal{B}_\pi$ by a partial function undefined on every argument.

For the inductive step, assume that the claim is true for l and for every σ-play π with $|\pi| = l_{\mathsf{max}} - l$ denote the $\mathcal{B}_\pi$-guide of σ_π by $\sigma^{\mathcal{B}_\pi}$. Given a σ-play π with $|\pi| = l_{\mathsf{max}} - l - 1$, there are three cases to consider:

- After π it is $\exists$ve who makes a move. Since π is a σ-play, σ provides a move $z = \sigma(\pi)$. Since πz is also a σ-play and $|\pi z| = l_{\mathsf{max}} - l$, by induction hypothesis there exists a winning strategy $\sigma_{\pi z}$ that is counter-independent in $\mathcal{B}_{\pi z}$. $\exists$ve can therefore win with the following strategy:

  $$\sigma_\pi(\rho) = \begin{cases} \sigma_{\pi z}(\rho) & \text{if } \pi z \text{ is a prefix of } \rho, \\ \sigma(\rho) & \text{otherwise.} \end{cases}$$

  Unless $\pi, \pi z \in \mathcal{B}$ and $z \in V$, the strategy σ_π is guided by $\sigma^{\mathcal{B}_\pi} = \sigma^{\mathcal{B}_{\pi z}}$ in $\mathcal{B}_\pi$. Otherwise it is guided by:

  $$\sigma^{\mathcal{B}_\pi}(\overline{v}) = \begin{cases} z & \text{if } \overline{v} = \mathsf{pos}(\pi), \\ \sigma^{\mathcal{B}_{\pi z}}(\overline{v}) & \text{otherwise.} \end{cases}$$

- After π $\forall$dam chooses a position v from a set $W \subseteq V$. For every such v, πv is a σ-play, $|\pi v| = l_{\mathsf{max}} - l$ and hence induction hypothesis provides $\sigma_{\pi v}$ guided by $\sigma^{\mathcal{B}_{\pi v}}$ in $\mathcal{B}_{\pi v}$. We combine strategies for all the possible choices from W:

  $$\sigma_\pi(\rho) = \begin{cases} \sigma_{\pi v}(\rho) & \text{if } \pi v \text{ is a prefix of } \rho, \\ \sigma(\rho) & \text{otherwise.} \end{cases}$$

  Such σ_π is guided in $\mathcal{B}_\pi$ by:

  $$\sigma^{\mathcal{B}_\pi}(\overline{v}) = \begin{cases} \sigma^{\mathcal{B}_{\pi v}}(\overline{v}) & \overline{v} \text{ has } \mathsf{pos}(\pi v) \text{ as a prefix,} \\ \text{undefined} & \text{otherwise.} \end{cases}$$

- After π $\forall$dam updates the current counters ctr to ctr'. The only interesting case is when the current rank r is nonstandard and so ctr' is given by a choice of an ordinal $\alpha < \mathsf{ctr}(r)$ (the case with standard r is similar to the first one). Denote such ctr' by ctr_α and $\pi\mathsf{ctr}_\alpha$ by π_α. For every $\alpha < \mathsf{ctr}(r)$ the play π_α is consistent with σ and $|\pi_\alpha| = l_{\mathsf{max}} - l$, so induction hypothesis gives us σ_{π_α} guided by $\sigma^{\mathcal{B}_{\pi_\alpha}}$ in $\mathcal{B}_{\pi_\alpha}$.

  Observe that for plays π_α, π_β leading to configurations γ_α and γ_β, respectively, we have $\gamma_\beta \preccurlyeq_{\mathsf{better}} \gamma_\alpha$ whenever $\alpha < \beta$. It follows that if after π $\forall$dam chooses α, $\exists$ve may as well continue as if he picked β. Denote such strategy by $\sigma_{[\beta/\alpha]}$.

  Importantly, if $\alpha \leq \beta$ and $\sigma^{\mathcal{B}_{\pi_\beta}}$ guides σ_{π_β} in $\mathcal{B}_{\pi_\beta}$, then it also guides $\sigma_{[\beta/\alpha]}$ in $\mathcal{B}_{\pi_\alpha}$. This is because whenever $\sigma_{[\beta/\alpha]}$-plays $\pi_\alpha\xi$ and $\pi_\alpha\xi v$

belong to $\mathcal{B}_{\pi_\alpha}$ and v is chosen by $\exists$ve, there exists $\pi_\beta\xi' \preccurlyeq_{\mathsf{better}} \pi_\alpha\xi$ such that $\pi_\beta\xi'$ and $\pi_\beta\xi' v$ are σ_{π_β}-plays. Since $\pi_\beta\xi' \preccurlyeq_{\mathsf{better}} \pi_\alpha\xi$ implies $\mathsf{pos}(\pi_\alpha\xi) = \mathsf{pos}(\pi_\beta\xi')$, by properness of $\mathcal{B}$ both $\pi_\beta\xi'$ and $\pi_\beta\xi' v$ belong to $\mathcal{B}_{\pi_\beta}$. Hence, $\sigma^{\mathcal{B}_{\pi_\beta}}(\mathsf{pos}(\pi_\alpha\xi)) = \sigma^{\mathcal{B}_{\pi_\beta}}(\mathsf{pos}(\pi_\beta\xi')) = v$, as desired.

There are two cases to consider, depending on whether $\mathsf{ctr}(r)$ is a limit ordinal or not. If it is a successor ordinal then there is a maximal α that can be chosen by $\forall$dam. In that case, $\exists$ve uses the strategy:

$$\sigma_\pi(\rho) = \begin{cases} \sigma_{[\alpha/\beta]}(\rho) & \text{if } \pi_\beta \text{ is a prefix of } \rho, \\ \sigma(\rho) & \text{otherwise,} \end{cases}$$

guided in $\mathcal{B}_\pi$ by $\sigma^{\mathcal{B}_\pi} = \sigma^{\mathcal{B}_{\pi_\alpha}}$.

On the other hand, if $\mathsf{ctr}(r)$ is a limit ordinal then there is no maximal α that $\forall$dam can choose, and for each of his choices $\exists$ve might have used a different σ_{π_α}. However, by assumption the set of positions that appear in $\mathcal{B}$ is finite. As a consequence, there are only finitely many possible partial functions guiding σ_{π_α} in $\mathcal{B}_{\pi_\alpha}$ and we may find $\sigma^{\mathcal{B}_\pi}$ such that σ_{π_α} is guided in $\mathcal{B}_{\pi_\alpha}$ by $\sigma^{\mathcal{B}_\pi}$ for arbitrarily big $\alpha < \mathsf{ctr}(r)$. Define:

$$\sigma_\pi(\rho) = \begin{cases} \sigma_{[\alpha'/\alpha]}(\rho) & \text{if } \pi_\alpha \text{ is a prefix of } \rho, \\ \sigma(\rho) & \text{otherwise,} \end{cases}$$

where $\alpha' \geq \alpha$ is the least number greater than α with $\sigma^{\mathcal{B}_{\pi_\alpha}} = \sigma^{\mathcal{B}_\pi}$. By design, σ_π is guided by $\sigma^{\mathcal{B}_\pi}$ in $\mathcal{B}_\pi$.

□

Positionality (and the lack thereof) is one of the aspects in which parity and countdown games differ. However, both types of games share the property that the players may use strategies that avoid unnecessary repetitions of positions. A *bad loop* for player P in a countdown play π is a suffix $v_i ... v_j$ of the underlying positions $\mathsf{pos}(\pi)$ such that $v_i = v_j$, the rank of v_i belongs to P and no v_k with $i < k < j$ has a more important rank.

Proposition 4.3.5. *Consider a countdown game $\mathcal{G} = (V, E, \mathsf{rank}, \mathsf{ctr}_I)$. If player P wins, then (s)he wins with a strategy σ that avoids bad loops, meaning that no σ-play contains a loop bad for P.*

Proof. The reasoning behind the above fact is essentially the same as for plain parity games. In a countdown game, if a play moves from a position v to itself via a path without visiting ranks higher than $\mathsf{rank}(v)$, then all the counters for lower ranks are reset and those for higher ranks remain unchanged. It follows that the resulting configuration is at least as good for the opponent P' of the owner P of $\mathsf{rank}(v)$ as the one at the previous visit to v. Hence, P' can repeat the strategy from that moment, and either eventually the game stops looping on v via lower ranks or P looses. This means that in order to win, P must have a strategy that avoids such loops, and therefore (s)he may use that strategy immediately. $\square$

4.4 Countdown Automata

As mentioned, countdown games are designed in a way that allows to lift the classical correspondence between logic, games and automata to the countdown setting. As with simple parity games and μ-ML, the bridge between countdown games and $\mu^{<\infty}$-ML is given by *countdown automata.* Such an automaton is almost the same as an ordinary parity automaton, except that the semantic game defining the language of a given automaton is a countdown game rather than a parity game. On the syntactic level, a countdown game is just a parity game extended with a set $\mathcal{D} \subseteq \mathcal{R}$ of nonstandard ranks and an initial counter assignment $\mathsf{ctr}_I : \mathcal{D} \to \mathsf{Ord}$. By analogy, we obtain a countdown automaton by extending a parity automaton with these two missing components.

Definition 4.4.1. A *countdown automaton* $\mathcal{A}$ is a tuple $(Q, q_I, \delta, \mathsf{rank}, \mathsf{ctr}_I)$ such that $(Q, q_I, \delta, \mathsf{rank})$ is a parity automaton and $\mathsf{ctr}_I : \mathcal{D} \to \mathsf{Ord}$ is an initial counter assignment for a subset $\mathcal{D} \subseteq \mathcal{R}$ of nonstandard ranks. Explicitly, $\mathcal{A}$ consists of:

- a finite set of states $Q = Q_\exists \sqcup Q_\forall$ divided between two players;
- an initial state $q_I \in Q$;
- a transition function:

$$\delta : Q \to \mathcal{P}(Q \sqcup \mathsf{Lit} \sqcup \mathsf{Var}) \sqcup (\mathsf{Act} \times Q)$$

- a rank function $\mathsf{rank} : Q \to \mathcal{R}$;
- a set $\mathcal{D} \subseteq \mathcal{R}$ of nonstandard ranks and an assignment of initial counter values $\mathsf{ctr}_I : \mathcal{D} \to \mathsf{Ord}$, as in a countdown game.

Analogously to parity automata, the language of a countdown automaton is given by a semantic game.

Definition 4.4.2. Fix an automaton $\mathcal{A} = (Q, q_I, \delta, \mathsf{rank}, \mathsf{ctr}_I)$, a model $\mathcal{M}$ and a valuation $\mathsf{val} : \mathsf{Var} \to \mathcal{P}(M)$. The semantic game $\mathcal{G}^{\mathsf{val}}(\mathcal{A})$ is the countdown game given as the extension of the semantic game induced by the *parity* automaton $(Q, q_I, \delta, \mathsf{rank})$ with $\mathcal{D} \subseteq \mathcal{R}$ and $\mathsf{ctr}_I : \mathcal{D} \to \mathsf{Ord}$. Explicitly, $\mathcal{G}^{\mathsf{val}}(\mathcal{A}) = (V, E, \mathsf{rank}', \mathsf{ctr}_I)$ where positions are of the form:

$$V = M \times (Q \sqcup \mathsf{Lit} \sqcup \mathsf{Var})$$

and the edge relation E is defined as with parity automata in Subsection 2.4.3. That is, in a position (m, q) for $q \in Q$:

- if $\delta(q) \subseteq Q \sqcup \mathsf{Lit} \sqcup \mathsf{Var}$, outgoing edges are:

$$\{((\mathsf{m}, q), (\mathsf{m}, z)) \mid z \in \delta(q)\},$$

- if $\delta(q) = (\mathsf{a}, p)$, outgoing edges are:

$$\{((\mathsf{m}, q), (\mathsf{n}, p)) \mid \mathsf{m} \xrightarrow{\mathsf{a}} \mathsf{n}\}.$$

There are no outgoing edges from positions $(\mathsf{m}, x), (\mathsf{m}, \tau)$ nor $(\mathsf{m}, \neg\tau)$ for $x \in \mathsf{Var}$ and $\tau \in \mathsf{Prop}$.

For $q \in Q$, the owner of the position (m, q) is the owner of the state q, and $\mathsf{rank}'(\mathsf{m}, q) = \mathsf{rank}(q)$. For $x \in \mathsf{Var}$, the position (m, x) belongs to $\forall$dam if $\mathsf{m} \in \mathsf{val}(x)$ and to $\exists$ve otherwise. Similarly, (m, τ) (or $(\mathsf{m}, \neg\tau)$) belongs to $\exists$ve iff $\mathsf{m} \notin \tau^{\mathcal{M}}$ (or $\mathsf{m} \in \tau^{\mathcal{M}}$, respectively). The rank' of (m, x), (m, τ) and $(\mathsf{m}, \neg\tau)$ can be set to an arbitrary standard rank so that it does not affect the outcome of the game.

The set $\mathcal{D} \subseteq \mathcal{R}$ of nonstandard ranks and the initial counter assignment ctr_I are taken from the automaton.

The semantics $[\![\mathcal{A}]\!]^{\mathsf{val}} \subseteq M$ of an automaton $\mathcal{A}$ is the set of all points $\mathsf{m} \in M$ for which the configuration $((\mathsf{m}, q_I), \mathsf{ctr}_I, \mathsf{cdn})$ in the game $\mathcal{G}^{\mathsf{val}}(\mathcal{A})$ is winning for $\exists$ve, in which case we say that $\mathcal{A}$ accepts $\mathcal{M}, \mathsf{m}$. The language of $\mathcal{A}$ is the class of all pointed models it accepts.

We will now explain the translations between logic and automata in turn.

4.4.1 From Formulae to Automata: Game Semantics

As with μ-ML, every countdown formula $\varphi \in \mu^{<\infty}$-$\mathsf{ML}$ can be viewed as a countdown automaton $\mathcal{A}_\varphi$ such that $[\![\varphi]\!]^{\mathsf{val}} = [\![\mathcal{A}_\varphi]\!]^{\mathsf{val}}$ for every model $\mathcal{M}$ and valuation val. The construction is very close to the one for plain μ-ML, with two key differences.

First, Example 4.2.5 demonstrates that an equality analogous to the Bekić principle (2.9) is not valid for countdown operators μ^α and ν^α, and Theorem 4.6.2 shows that in fact there is no other way around, as vectorial formulae have more expressive power. This requires an adaptation of the game to the more general vectorial setting. Instead of having a variable x and a formula θ such that $\eta x.\theta$ binds x, we now have *tuples* $\overline{\theta} = \langle\theta_1, ..., \theta_d\rangle$ and $\overline{x} = \langle x_1, ..., x_d\rangle$ such that $\eta^\alpha \overline{x}.\overline{\theta}$ binds $\overline{x}$. Then, instead of the deterministic move from x to θ as in the classical case, the vectorial semantic game moves from x_i to the formula θ_i with a matching coordinate $1 \leq i \leq d$.

Second, $\mu^{<\infty}$-ML features countdown operators, which define languages beyond the regular ones. Since our design goal was to lift the correspondence between logic and games to the richer countdown setting, we now want to define a *countdown*, rather than a *simple parity* game, in hope it will capture the semantics of $\mu^{<\infty}$-ML. In order to turn a simple parity game into a countdown one, it suffices to determine a subset $\mathcal{D} \subseteq \mathcal{R}$ of nonstandard ranks together with an assignment $\mathsf{ctr}_I : \mathcal{D} \to \mathsf{Ord}$ of initial counter values. To this end, we define the irrelevant lowest rank 0 to be standard and for all the other ranks, i.e. ranks r of immediate subformulae $\theta_1, ..., \theta_d$ of a fixpoint formula $\eta_i^\alpha \overline{x}.\overline{\theta}$, we look at the index $\alpha \in \mathsf{Ord}_\infty$. If $\alpha = \infty$, then η_i^α is a usual fixpoint operator and hence we leave its rank standard. Otherwise, we put $r \in \mathcal{D}$ with $\mathsf{ctr}_I(r) = \alpha$.

Applying the above modifications, we obtain the following definition of a countdown automaton $\mathcal{A}_\varphi = (Q, q_I, \delta, \mathsf{rank}, \mathsf{ctr}_I)$ corresponding to a given formula φ:

- the set Q of states and its ownership is the same as in the simple parity case:
$$Q = \mathsf{SubFor}(\varphi) - (\mathsf{Lit} \sqcup \mathsf{FreeVar}(\varphi))$$
$\exists$ve owns formulae with $\vee$ and $\langle \mathsf{a} \rangle$ as the topmost connective and $\forall$dam these with $\wedge$ and $[\mathsf{a}]$; ownership of fixpoint subformulae and countdown subformulae is irrelevant;

- $q_I = \varphi$;

- the transition function is defined by cases:
  - $\delta(\theta_1 \vee \theta_2) = \delta(\theta_1 \wedge \theta_2) = \{\theta_1, \theta_2\}$,
  - $\delta(\langle \mathsf{a} \rangle \theta) = \delta([\mathsf{a}]\theta) = (\mathsf{a}, \theta)$,
  - $\delta(\eta_i^\alpha \overline{x}.\overline{\theta}) = \{\theta_i\}$ (for $\eta = \mu$ or $\eta = \nu$),
  - $\delta(x) = \{\theta_i\}$, where $\eta_j^\alpha(x_1, ..., x_d).(\theta_1, ..., \theta_d)$ is the (unique) subformula of φ binding x with $x = x_i$.

- For the ranking function, assume that the lowest rank in $\mathcal{R}$ is standard and call it 0 (ownership of this rank does not matter). Then let rank assign 0 to all subformulae of φ except for immediate subformulae of fixpoint operators. To those, assign ranks in such a way that subformulae have strictly smaller ranks than their superformulae, and for every subformula $\eta_i^\alpha \overline{x}.\overline{\varphi}$:
  - all formulae in the tuple $\overline{\varphi}$ have the same rank r,
  - r belongs to $\exists$ve if $\eta = \mu$ and to $\forall$dam if $\eta = \nu$, and
  - if $\alpha = \infty$ then r is standard, otherwise it is nonstandard and $\mathsf{ctr}_I(r) = \alpha$.

We put $\mathcal{G}^{\mathsf{val}}(\varphi) = \mathcal{G}^{\mathsf{val}}(\mathcal{A}_\varphi)$.

Theorem 4.4.3 (Countdown Adequacy)**.** *For every model $\mathcal{M}$ and valuation* val*,* $[\![\varphi]\!]^{\mathsf{val}} = [\![\mathcal{A}_\varphi]\!]^{\mathsf{val}}$*.*

Proof. Unfolding the definition of $[\![\mathcal{A}_\varphi]\!]^{\mathsf{val}}$ from Definition 4.4.2 we prove that:

$$\mathsf{m} \in [\![\varphi]\!]^{\mathsf{val}} \iff \exists\text{ve wins } \mathcal{G}^{\mathsf{val}}(\varphi) \text{ from } ((\mathsf{m}, \varphi), \mathsf{ctr}_I, \mathsf{cdn}) \tag{4.5}$$

for every $\mathsf{m} \in M$ and valuation val. As with the Adequacy Theorem 2.4.7 for the classical μ-calculus, the proof proceeds by induction on the complexity of the formula φ. The only new cases are $\varphi = \mu_i^\alpha \overline{x}.\overline{\psi}$ and $\varphi = \nu_i^\alpha \overline{x}.\overline{\psi}$ for $\overline{x} = \langle x_1, ..., x_d \rangle$, $\overline{\varphi} = \langle \varphi_1, ..., \varphi_d \rangle$ and $\alpha \in \mathsf{Ord}_\infty$. Since the cases with μ and ν are symmetric we only consider the first one. Moreover, we focus on the case with $\alpha \in \mathsf{Ord}$ as in the remaining one $\alpha = \infty$ the proof only simplifies. By definition of the semantics, $[\![\mu_i^\alpha \overline{x}.\overline{\psi}]\!]^{\mathsf{val}} = \pi_i(F_\mu^\alpha)$ for:

$$F(H_1, ..., H_d) = ([\![\psi_1]\!]^{\mathsf{val}'}, ..., [\![\psi_d]\!]^{\mathsf{val}'})$$

where $\mathsf{val}' = \mathsf{val}[x_1 \mapsto H_1, ..., x_d \mapsto H_d]$. The proof of (4.5) is similar to the analogous equivalence (2.10) for plain μ-ML. The first difference is that we deal with a countdown, rather than a fixpoint operator. For that reason, we apply Theorem 4.1.2 in place of Theorem 2.2.8. The second difference is the presence of vectorial formulae. Because of this, we use the natural isomorphism:

$$\iota : \mathcal{P}(M)^d \cong \mathcal{P}(M \times d)$$

which maps every tuple of sets $(H_1, ..., H_d) \in \mathcal{P}(H)^d$ to the set $\{(\mathsf{n}, i) \mid \mathsf{n} \in H_i\} \in \mathcal{P}(M \times \{1, ..., d\})$. Such ι allows us to view F as an operation $\widehat{F} = \iota \circ F \circ \iota^{-1}$ on $\mathcal{P}(M \times \{1, ..., d\})$. Then:

$$\begin{aligned} \mathsf{m} \in [\![\varphi]\!]^{\mathsf{val}} &\iff \mathsf{m} \in \pi_i(F_\mu^\alpha) \\ &\iff (\mathsf{m}, i) \in \widehat{F}_\mu^\alpha \\ &\iff \exists\text{ve wins } \mathcal{G}_\mu^{<\infty}(\widehat{F}) \text{ from } ((\mathsf{m}, i), \alpha, \mathsf{cdn}). \end{aligned}$$

The first equivalence is the definition of the meaning of μ^α. The second one follows from equality $\iota(F_\mu^\beta) = \widehat{F}_\mu^\beta$ proven by straightforward induction on $\beta \in \mathsf{Ord}$. The third one follows from Theorem 4.1.2. Hence, in order to prove (4.5) it suffices to show:

$$\mathcal{G}, ((\mathsf{m}, i), \alpha, \mathsf{cdn}) \bowtie \mathcal{G}', ((\mathsf{m}, \varphi), \mathsf{ctr}_I, \mathsf{cdn})$$

where:

$$\mathcal{G} = \mathcal{G}_\mu^{<\infty}(\widehat{F}) \qquad \text{and} \qquad \mathcal{G}' = \mathcal{G}^{\mathsf{val}}(\varphi).$$

Since the game on the right begins with a vacuous counter update (as φ has an irrelevant rank) followed by a deterministic ε-transition to ψ_i, the crux of the proof is to show:

$$\mathcal{G}, ((\mathsf{m}, i), \alpha, \mathsf{cdn}) \bowtie \mathcal{G}', ((\mathsf{m}, \psi_i), \mathsf{ctr}_I, \mathsf{cdn}). \tag{4.6}$$

We will use the Decomposition Lemma 2.2.5 with $S = \emptyset$ and sets of stopping configurations:

$$\mathcal{S}_+ = (M \times \{1, ..., d\}) \times \mathsf{Ord} \times \{\mathsf{cdn}\} \qquad \text{and} \qquad \mathcal{S}'_+ = (M \times \{\psi_1, ..., \psi_d\}) \times \mathsf{Ord}^{\mathcal{D}} \times \{\mathsf{cdn}\}.$$

$\mathcal{S}_+$ and $\mathcal{S}'_+$ are precisely the sets of configurations with the most important ranks r and r' of $\mathcal{G}$ and $\mathcal{G}'$, respectively. Consider the relation $S_+ \subseteq \mathcal{S}_+ \times \mathcal{S}'_+$ defined as:

$$((\mathsf{n}, j), \beta, \mathsf{cdn}) \ S_+ \ ((\mathsf{n}', \psi_{j'}), \mathsf{ctr}, \mathsf{cdn})$$
$$\iff$$
$$\mathsf{n} = \mathsf{n}',\ j = j' \text{ and } \mathsf{ctr}(r') = \beta$$

for all $((\mathsf{n}, j), \beta, \mathsf{cdn}) \in \mathcal{S}_+$ and $((\mathsf{n}', \psi_{j'}), \mathsf{ctr}, \mathsf{cdn}) \in \mathcal{S}'_+$. With the above definitions we get (4.6) from the lemma once we prove:

$$\mathcal{G}, ((\mathsf{n}, j), \beta, \mathsf{cdn}) \bowtie_{S_+} \mathcal{G}', ((\mathsf{n}, \psi_j), \mathsf{ctr}, \mathsf{cdn}) \tag{4.7}$$

for all $\mathsf{n} \in M$, $\beta \in \mathsf{Ord}$, $j \leq d$ and ctr such that $\mathsf{ctr}(r') = \beta$. Observe that both games start with $\exists$ve, the owner of r and r', choosing some $\kappa < \beta$. This κ becomes the new value of the only counter and in $\mathcal{G}$ and the value of the counter for r' in $\mathcal{G}'$ (and all the remaining counters get reset, as r' is the most important). Therefore, it suffices to prove:

$$\mathcal{G}, ((\mathsf{n}, j), \kappa, \mathsf{psn}) \bowtie_{S_+} \mathcal{G}', ((\mathsf{n}, \psi_j), \mathsf{ctr}_I[r' \mapsto \kappa], \mathsf{psn})$$

for all $\kappa < \beta$. To that end take arbitrary $H_1, ..., H_d \subseteq M$ and denote by $\widehat{H} = \iota(\overline{H})$ the result of applying the isomorphism ι to the tuple $\overline{H}$. We complete the proof by showing that the following are equivalent:

1. $\exists$ve has a non-loosing strategy σ for $\mathcal{G}|\mathcal{S}_+, ((\mathsf{n}, j), \kappa, \mathsf{psn})$ with exit configurations $\mathsf{exit}(\sigma) \subseteq \mathcal{S}_+$ equal:
$$\widehat{H} \times \{\kappa\} \times \{\mathsf{cdn}\},$$
2. $(\mathsf{n}, j) \in \widehat{F}(\widehat{H})$,
3. $\mathsf{n} \in \pi_j(F(\overline{H}))$,
4. $\exists$ve has a winning strategy for $\mathcal{G}^{\mathsf{val}'}(\psi_j), (\mathsf{n}, \psi_j)$ where $\mathsf{val}' = \mathsf{val}[x_1 \mapsto H_1, ..., x_d \mapsto H_d]$,
5. $\exists$ve has a non-loosing strategy σ' for $\mathcal{G}'|\mathcal{S}'_+, ((\mathsf{n}, \psi_j), \mathsf{ctr}_I[r' \mapsto \kappa], \mathsf{psn})$ with exit configurations $\mathsf{exit}(\sigma') \subseteq \mathcal{S}'_+$ included in:
$$\{((\mathsf{n}', \psi_k), \mathsf{ctr}, \mathsf{cdn}) \mid k \leq d, \mathsf{n}' \in H_k, \mathsf{ctr}(r') = \kappa\}.$$

Equivalence (1) $\iff$ (2) follows from the definition of $\mathcal{G}$; (2) $\iff$ (3) from the definition of ι and $\widehat{F}$ and (3) $\iff$ (4) from the definition of F and the induction hypothesis (4.5) applied to ψ_j. The last equivalence (4) $\iff$ (5) holds because the two games $\mathcal{G}^{\mathsf{val}'}(\psi_j), (\mathsf{n}, \psi_j)$ and $\mathcal{G}'|\mathcal{S}'_+, ((\mathsf{n}, \psi_j), \mathsf{ctr}_I[r' \mapsto \kappa], \mathsf{psn})$ are isomorphic until a move to position (n', x_k) for some $k \leq d$ and $\mathsf{n}' \in M$ (in particular, in $\mathcal{G}'|\mathcal{S}'_+$ the value $\mathsf{ctr}(r') = \kappa$ cannot be changed because r' cannot be visited before the game stops). If the games ever reach such (n', x_k), in $\mathcal{G}^{\mathsf{val}'}(\psi_j)$ $\exists$ve wins iff $\mathsf{n}' \in H_k$ whereas $\mathcal{G}'$ moves deterministically to (n', ψ_k) and stops there. This establishes equivalence between winning strategies in (3) and non-loosing strategies with appropriate exit-configurations in (4). $\square$

Example 4.4.4. For $\mathsf{Act} = \{\mathsf{a}\}$, consider the formula $\varphi = \mu^\omega x.\Box x$ from Example 4.2.1. The automaton $\mathcal{A}_\varphi$ has three states: $Q = \{\varphi, \Box x, x\}$, with φ the initial state, and the transition function comprises two deterministic ε-transitions and one modal transition:

$$\delta(\varphi) = \{\Box x\}, \qquad \delta(\Box x) = (\mathsf{a}, x), \qquad \delta(x) = \{\Box x\}.$$

The state $\Box x$ is owned by $\forall$dam; ownership of the other two states does not matter. The automaton uses two ranks, $0 < 1$, where 0 is standard and 1 is nonstandard, assigned to states by: $\mathsf{rank}(\varphi) = \mathsf{rank}(x) = 0$ and $\mathsf{rank}(\Box x) = 1$. Rank 1 is owned by $\exists$ve; ownership of rank 0 does not matter. (Note how the state $\Box x$ is owned by $\forall$dam, but its rank is owned by $\exists$ve). The initial counter value is $\mathsf{ctr}_I(1) = \omega$.

Now consider any model $\mathcal{M}$. Since Act has only one element, $\mathcal{M}$ is simply a directed graph $(M, \to)$. The semantic game $\mathcal{G}(\varphi)$ on $\mathcal{M}$ (φ has no free variables, so neither has $\mathcal{A}_\varphi$ and we need not consider valuations val) has positions of the form (m, q) where $\mathsf{m} \in M$ and $q \in Q$, with ownership and rank inherited from q. Edges are of the form:

- $((\mathsf{m}, \varphi), (\mathsf{m}, \Box x))$ and $((\mathsf{m}, x), (\mathsf{m}, \Box x))$ – the ε-edges,
- $((\mathsf{m}, \Box x), (\mathsf{n}, x))$ such that $\mathsf{m} \to \mathsf{n}$ is an edge in $\mathcal{M}$ – the modal edges.

Configurations of the game arise from positions together with counter assignments; there is only one nonstandard rank, so a counter assignment is simply an ordinal.

For a point $\mathsf{m} \in \mathcal{M}$, the default initial configuration of the game is the countdown configuration $((\mathsf{m}, \varphi), \omega, \mathsf{cdn})$. A play that begins in this configuration proceeds as follows:

1. The first two moves are deterministic: first to the positional configuration $((\mathsf{m}, \varphi), \omega, \mathsf{cdn})$ (since the rank 0 of φ is standard and smaller than 1); and then to the countdown configuration $((\mathsf{m}, \Box x), \omega, \mathsf{cdn})$.

2. $\exists$ve, as the owner of the rank 1 of $\Box x$, makes the next move: she chooses a number $k < \omega$, and the games moves to the positional configuration $((\mathsf{m}, \Box x), k, \mathsf{psn})$.

3. $\forall$dam owns the position, so he makes the next move: he chooses a point $\mathsf{n} \in M$ such that $\mathsf{m} \to \mathsf{n}$, and the game moves to the countdown configuration $((\mathsf{n}, x), k, \mathsf{cdn})$.

4. The rank 0 of x is standard and smaller than 1, so in the next move the counter does not change and the game moves to $((\mathsf{n}, x), k, \mathsf{psn})$. The next move is also deterministic, to the countdown configuration $((\mathsf{n}, \Box x), k, \mathsf{cdn})$. The game then goes back to step 2. above, with k in place of ω.

From this it is clear that $\exists$ve wins from $((\mathsf{m}, \varphi), \omega, \mathsf{cdn})$ if and only if there is a finite bound on the lengths of paths starting in m, as stated in Example 4.2.1.

4.4.2 From Automata to Formulae

We provide a theorem that directly generalizes Theorem 2.4.8 to the countdown setting.

Theorem 4.4.5. *For every countdown automaton $\mathcal{A}$ there exists a formula $\varphi_{\mathcal{A}}$ of $\mu^{<\infty}$-ML such that $[\![\mathcal{A}]\!]^{\mathsf{val}} = [\![\varphi_{\mathcal{A}}]\!]^{\mathsf{val}}$ for every model $\mathcal{M}$ and valuation* val.

Proof. Syntactically, the proof is very similar to the classical one. Fix a countdown automaton $\mathcal{A} = (Q, q_I, \delta, \mathsf{rank}, \mathsf{ctr}_I)$. As with parity automata, we only consider the case when $\mathcal{A}$ has no free variables, the highest rank r_{max} is not assigned to any state and every other rank is assigned to at least one state. Note that, unlike in the classical setting, we *cannot* assume that each rank is assigned to at most one state, though. For convenience, we assume that all the initial counter values are nonzero (otherwise the owner P of the corresponding rank r looses immediately upon entering a state with rank r, so we may turn all the states with that rank into dead-end states owned by P and having the standard rank 0). Taking variables $\mathsf{Var}_Q = \{x_q \mid q \in Q\}$ with distinct x_q for every $q \in Q$ denote:

$$Q_{r\leq} = \{q \in Q \mid r \leq \mathsf{rank}(q)\} \qquad \text{and} \qquad \mathsf{Var}_{r\leq} = \{x_q \mid q \in Q_{r\leq}\}.$$

Proceeding by induction on r, for each state q we construct a formula $\psi_{r,q}$ over Var_Q with all free variables in $\mathsf{Var}_{r\leq}$ and all bound variables outside of $\mathsf{Var}_{r\leq}$. We simplify notation: for avery free variable x_p in $\psi_{r,q}$ we identify all the occurrences of x_p as subformulae of $\psi_{r,q}$. For syntactic simplicity, we assume that all the constructed formulae, when seen as automata, have the same ranks $\mathcal{R}'$, nonstandard ranks $\mathcal{D}'$ and initial counter assignment

$\mathsf{ctr}'_I \in \mathsf{Ord}^{\mathcal{D}'}$ (the rank function $\mathsf{rank}_{r,q} : \mathsf{SubFor}(\psi_{r,q}) \to \mathcal{R}'$ may depend on r and q and need not be surjective).

The goal of our construction is to have:

$$\mathcal{G}(\mathcal{A}), ((\mathsf{m}, q), \mathsf{ctr}_I, \mathsf{cdn}) \bowtie_{S_{r\leq}} \mathcal{G}(\psi_{r,q}), ((\mathsf{m}, \psi_{r,q}), \mathsf{ctr}'_I, \mathsf{cdn}), \tag{4.8}$$

for every $\mathsf{m} \in \mathcal{M}$ where $S_{r\leq}$ is the relation:

$$S_{r\leq} = \{(((\mathsf{n}, p), \mathsf{ctr}, \mathsf{cdn}), ((\mathsf{n}, x_p), \mathsf{ctr}', \mathsf{cdn})) \mid p \in Q_{r\leq}\}$$

with domain and codomain:

$$\mathcal{S}_{r\leq} = M \times Q_{r\leq} \times \mathsf{Ord}^{\mathcal{D}} \times \{\mathsf{cdn}\} \qquad \text{and} \qquad \mathcal{S}'_{r\leq} = M \times \mathsf{Var}_{r\leq} \times \mathsf{Ord}^{\mathcal{D}'} \times \{\mathsf{cdn}\}.$$

The above definitions are analogous to the ones for the classical case (2.12), with sets $\mathcal{S}_{r\leq}$ and $\mathcal{S}_{r\leq}$ extended to countdown configurations (and not just mere positions) and the relation $S_{r\leq}$ adapted accordingly. Note that although we will refer to the counter values at some point in the proof, the relation $S_{r\leq}$ in the induction hypothesis ignores them.

The Base Case. The base case $r = 0$ is no different from the classical variant:

- if $\delta(q) = (\mathsf{a}, p)$:

$$\psi_{0,q} = \begin{cases} \langle \mathsf{a} \rangle p & \text{if } q \text{ belongs to } \exists\text{ve} \\ [\mathsf{a}]p & \text{if } q \text{ belongs to } \forall\text{dam} \end{cases}$$

- if $\delta(q) \subseteq Q \sqcup \mathsf{Lit} \sqcup \mathsf{Var}$:

$$\psi_{0,q} = \begin{cases} \bigvee \delta(q) & \text{if } q \text{ belongs to } \exists\text{ve} \\ \bigwedge \delta(q) & \text{if } q \text{ belongs to } \forall\text{dam.} \end{cases}$$

The game $\mathcal{G}(\mathcal{A})|\mathcal{S}_{0\leq}$ is the same as in the classical case, except that before stopping or ending it goes through a counter update. Such update is irrelevant, though, because we only care about the position in which the game stops and since the initial counter value is nonzero it can always be decremented.

The Inductive Step. For the inductive step, assume (4.8) for rank r, which gives us $\psi_{r,q}$ for each $q \in Q$ and denote the next rank by $r + 1$. We want to construct $\psi_{r+1,q}$ satisfying:

$$\mathcal{G}(\mathcal{A}), ((\mathsf{m}, q), \mathsf{ctr}_I, \mathsf{cdn}) \bowtie_{S_{r+1\leq}} \mathcal{G}(\psi_{r+1,q}), ((\mathsf{m}, \psi_{r+1,q}), \mathsf{ctr}'_I, \mathsf{cdn}) \tag{4.9}$$

for all $\mathsf{m} \in \mathcal{M}$. Let $p_1, ..., p_d$ be the states in Q that have rank r. We will use the decomposition:

$$S_{r\leq} = S_r \sqcup S_{r+1\leq} \tag{4.10}$$

where $S_r = S_{r\leq} - S_{r+1\leq}$ can be described explicitly as:

$$((\mathsf{n}, s), \mathsf{ctr}, \mathsf{cdn})\ S_r\ ((\mathsf{n}', \theta), \mathsf{ctr}', \mathsf{cdn})$$
$$\Longleftrightarrow$$
$$\mathsf{n} = \mathsf{n}' \text{ and there is } i \leq d \text{ such that } s = p_i \text{ and } \theta = x_{p_i}.$$

We introduce the same notation as in the classical case:

$$S = S_{r+1\leq} \subseteq \mathcal{S} \times \mathcal{S}' \qquad \text{where} \qquad \mathcal{S} = \mathcal{S}_{r+1\leq} \qquad \text{and} \qquad \mathcal{S}' = \mathcal{S}'_{r+1\leq}.$$

The Case with $q \in \{p_1, ..., p_d\}$. We start the construction with the case $q = p_i$ for some $i \leq d$. Taking formulae $\overline{\psi_{r,p}} = \langle \psi_{r,p_1}, ..., \psi_{r,p_d} \rangle$ constructed in the previous case and the corresponding variables $\overline{x_p} = \langle x_{p_1}, ..., x_{p_d} \rangle$ put:

$$\psi_{r+1,p_i} = \eta_i^\alpha \overline{x_p}.\overline{\psi_{r,p}}$$

where $\eta = \mu$ if r belongs to $\exists$ve and $\eta = \nu$ if r belongs to $\forall$dam with $\alpha = \mathsf{ctr}_I(r)$ if r is nonstandard and $\alpha = \infty$ otherwise. That is, we generalize the classical scalar construction to tuples $\overline{\psi_{r,p}}$ and $\overline{x_p}$ and add two indices: i denoting the coordinate of p_i among $p_1, .., p_d$ and α denoting the initial value of a counter corresponding to r (or ∞ if r is standard). The idea is similar to the classical case except that now at each visit to the rank r, which corresponds to unravelling of the countdown operator, its owner has to decrement the counter, starting from α. Formally, we denote the rank of all $\psi_{r,p_1}, ..., \psi_{r,p_d}$ by r' and prove:

$$\mathcal{G}(\mathcal{A}), ((\mathsf{m}, p_i), \mathsf{ctr}, \mathsf{cdn}) \bowtie_S \mathcal{G}(\psi_{r+1,p_i}), ((\mathsf{m}, \psi_{r,p}), \mathsf{ctr}', \mathsf{cdn})) \tag{4.11}$$

for all $\mathsf{m} \in M$ and $\mathsf{ctr}, \mathsf{ctr}'$ with $\mathsf{ctr}(r) = \mathsf{ctr}'(r')$. This suffices to prove our induction goal (4.9), as the game $\mathcal{G}(\psi_{r+1,p_i}), ((\mathsf{m}, \psi_{r+1,p_i}), \mathsf{ctr}'_I, \mathsf{cdn})$ starts with a pair of deterministic moves (a vacuous counter update and an ε-move) leading to $((\mathsf{m}, \psi_{r,p_i}), \mathsf{ctr}'_I, \mathsf{cdn})$ and equality $\mathsf{ctr}_I(r) = \mathsf{ctr}'_I(r')$ follows directly from the definitions. Consider sets of configurations:

$$\begin{aligned}
\mathcal{S}_+ &= M \times \{p_j \mid j \leq d\} \times \mathsf{Ord}^{\mathcal{D}} \times \{\mathsf{cdn}\} \\
\mathcal{S}_+^\circ &= M \times \{x_{p_j} \mid j \leq d\} \times \mathsf{Ord}^{\mathcal{D}'} \times \{\mathsf{cdn}\} \\
\mathcal{S}'_+ &= M \times \{\psi_{r,p_j} \mid j \leq d\} \times \mathsf{Ord}^{\mathcal{D}'} \times \{\mathsf{cdn}\}.
\end{aligned}$$

We will use relations $S_+ \subseteq S_+^\circ \subseteq \mathcal{S}_+ \times \mathcal{S}'_+$ and $R \subseteq S_+^\circ \times S'_+$ defined for every $\gamma = ((\mathsf{n}, p_j), \mathsf{ctr}, \mathsf{cdn}) \in \mathcal{S}_+$, $\gamma^\circ = ((\mathsf{n}^\circ, x_{p_{j^\circ}}), \mathsf{ctr}^\circ, \mathsf{cdn})$ and $\gamma' = ((\mathsf{n}', \psi_{r,p_{j'}}), \mathsf{ctr}', \mathsf{cdn}) \in \mathcal{S}'_+$ as follows:

$$\begin{aligned}
\gamma S_+^- \gamma' &\iff \mathsf{n} = \mathsf{n}' \text{ and } j = j', \\
\gamma^\circ R \gamma' &\iff \mathsf{n}^\circ = \mathsf{n}', \text{ and } j^\circ = j', \\
\gamma S_+ \gamma' &\iff \mathsf{n} = \mathsf{n}', j = j' \text{ and } \mathsf{ctr}(r) = \mathsf{ctr}'(r').
\end{aligned}$$

That is, R links configurations with equal points and the variable x_{p_j} matching the formula ψ_{r,p_j}; similarly with S_+^- except that it requires the state p_j to match with the formula ψ_{r,p_j}. Only S_+ looks at the counters: is strengthens S_+^- with an additional requirement that the counter values for r and r' are equal.

We will use the Decomposition Lemma 2.2.5 substituting $\mathcal{G} = \mathcal{G}(\mathcal{A})$ and $\mathcal{G}' = \mathcal{G}(\psi_{r+1,p_i})$. The sets $\mathcal{S}_+$ and $\mathcal{S}'_+$ only contain configurations with the most important rank r in $\mathsf{Conf} - \mathcal{S}$ and r' in $\mathsf{Conf}' - \mathcal{S}'$, respectively. Hence, the lemma implies (4.11) once we prove:

$$\mathcal{G}(\mathcal{A}), \gamma \bowtie_{S_+ \sqcup S} \mathcal{G}(\psi_{r+1,p_i}), \gamma' \tag{4.12}$$

for all $(\gamma, \gamma') \in S_+$. Fix such a pair with $\gamma = ((\mathsf{n}, p_j), \mathsf{ctr}, \mathsf{cdn})$ and $\gamma' = ((\mathsf{n}, \psi_{r,p_j}), \mathsf{ctr}', \mathsf{cdn})$ such that $\mathsf{ctr}(r) = \mathsf{ctr}'(r')$. We first prove a weaker claim with the special case $\mathsf{ctr} = \mathsf{ctr}_I$ and $\mathsf{ctr}' = \mathsf{ctr}'_I$ and S_+^- in place of S_+:

$$\mathcal{G}(\mathcal{A}), ((\mathsf{n}, p_j), \mathsf{ctr}_I, \mathsf{cdn}) \bowtie_{S_+^- \sqcup S} \mathcal{G}(\psi_{r+1,p_j}), ((\mathsf{n}, \psi_{r,p_j}), \mathsf{ctr}'_I, \mathsf{cdn}). \tag{4.13}$$

For that compose:

$$\begin{aligned} \mathcal{G}(\mathcal{A}), ((\mathsf{n}, p_j), \mathsf{ctr}_I, \mathsf{cdn}) &\bowtie_{S_{r\le}} \mathcal{G}(\psi_{r,p_j}), ((\mathsf{n}, \psi_{r,p_j}), \mathsf{ctr}'_I, \mathsf{cdn}) \\ &\bowtie_{R \sqcup \mathsf{Id}(\mathcal{S}')} \mathcal{G}(\psi_{r+1,p_j}), ((\mathsf{n}, \psi_{r,p_j}), \mathsf{ctr}'_I, \mathsf{cdn}) \end{aligned}$$

The first equivalence is the induction hypothesis (4.8) applied to ψ_{r,p_j}. The second one is true for the same reasons as in the analogous proof for μ-ML. The partial games:

$$\mathcal{G}(\psi_{r,p_j}) | \mathcal{S}_+^\circ \sqcup \mathcal{S}', ((\mathsf{n}, \psi_{r,p_j}), \mathsf{ctr}'_I, \mathsf{cdn})$$

and:

$$\mathcal{G}(\psi_{r+1,p_j}) | \mathcal{S}'_+ \sqcup \mathcal{S}', ((\mathsf{n}, \psi_{r,p_j}), \mathsf{ctr}'_I, \mathsf{cdn})$$

are isomorphic until they move to some $((\mathsf{n}', x_{p_k}), \mathsf{ctr}', \mathsf{cdn})$ in which case the first game stops and the second one fires a vacuous counter update, moves deterministically to $((\mathsf{n}', \psi_{r,p_k}), \mathsf{ctr}', \mathsf{cdn})$ and stops as well. Triviality of the update follows from the observation that none of the formulae we construct is equal to a variable (assuming the convention that $\bigvee\{\theta\} = \theta \vee \theta$ and symmetrically for $\bigwedge$), so no variable is an immediate subformula of a countdown operator and hence all variables have rank 0. Since:

$$\begin{aligned} S_{r\le} \circ (R \sqcup \mathsf{Id}(\mathcal{S}')) &= (S_r \sqcup S) \circ (R \sqcup \mathsf{Id}(\mathcal{S}')) \\ &= (S_r \circ R) \sqcup (S \circ \mathsf{Id}(\mathcal{S}')) \\ &= S_+^- \sqcup S \end{aligned}$$

the two equivalences compose into (4.13). To get (4.12) we need to strengthen (4.13) in two ways:

(i) replace ctr_I and ctr'_I with any pair of counter assignments ctr and ctr' using the assumption $\mathsf{ctr}(r) = \mathsf{ctr}'(r')$ and

(ii) replace S_+^- with S_+, i.e. impose an additional requirement that if the games stop in some $(((\mathsf{n}', p_k), \mathsf{ctr}'', \mathsf{cdn}), ((\mathsf{n}', \psi_{r,p_k}), \mathsf{ctr}''', \mathsf{cdn})) \in S_+^-$ then actually $\mathsf{ctr}''(r) = \mathsf{ctr}'''(r')$.

For the first item (i) observe that both games start with the owner of r and r' (which is the same player) choosing $\beta < \mathsf{ctr}_I(r) = \mathsf{ctr}'_I(r')$ and resetting counters for all the less important ranks. These ranks include the only ones that can be reached before the games stop: ranks smaller than r in $\mathcal{G}(\mathcal{A})$ and ranks of strict subformulae of ψ_{r+1,p_j} in $\mathcal{G}(\psi_{r+1,p_j})$. Since after this initial decrement counters for all the reachable ranks have initial values, it follows that instead of ctr_I and ctr'_I we could take any ctr and ctr' provided that $0 < \mathsf{ctr}(r)$ iff $0 < \mathsf{ctr}'(r')$ (which follows from $\mathsf{ctr}_I(r) = \mathsf{ctr}'_I(r')$). Moreover, after this first countdown step the counters for r and r' will never be decremented again before the games stop. Since the equality $\mathsf{ctr}_I(r) = \mathsf{ctr}'_I(r')$ implies that the owner of r and r' can choose the same value at the beginning of both games, this gives us the second item (ii).

The Case with $q \notin \{p_1, ..., p_d\}$. We are left with the construction of $\psi_{r+1,q}$ for $q \notin \{p_1, ..., p_d\}$. Put:

$$\psi_{r+1,q} = \psi_{r,q}[x_{p_1} \mapsto \psi_{r+1,p_1}, ..., x_{p_d} \mapsto \psi_{r+1,p_d}].$$

That is, we take $\psi_{r,q}$ and replace every x_{p_i} with the respective ψ_{r+1,p_i} which we have just constructed.

Let $\theta_1, ...\theta_l$ be all the subformulae of $\psi_{r+1,q}$ which are copies of one of $\psi_{r+1,p_1}, ..., \psi_{r+1,p_d}$. We denote the subformula of θ_k that is a copy of $\psi \in \mathsf{SubFor}(\psi_{r+1,p})$ by ψ^k. Once the game reaches a particular ψ^k_{r,p_i}, the only reachable formulae are subformulae of θ_k. This is because a move leaving $\mathsf{SubFor}(\theta_k)$ would require existence of a variable x_s free in θ_k (meaning that $x_s \in \mathsf{Var}_{r+1\leq}$) and bound in $\psi_{r,q}$ (so $x_s \notin \mathsf{Var}_{r\leq}$) which is a contradiction. Thanks to this, we may assume that all the formulae $\{\psi^k_{r,p_i} \mid i \leq d, k \leq l\}$ have the same rank independent of k and denote it by r'.

We use sets of configurations and relations almost the same as in the previous case except that we deal with several copies of every ψ_{r,p_j}:

$$\begin{aligned}
\mathcal{S}_+ &= M \times \{p_j \mid j \leq d\} \times \mathsf{Ord}^{\mathcal{D}} \times \{\mathsf{cdn}\} \\
\mathcal{S}^\circ_+ &= M \times \{x_{p_j} \mid j \leq d\} \times \mathsf{Ord}^{\mathcal{D}'} \times \{\mathsf{cdn}\} \\
\mathcal{S}'_+ &= M \times \{\psi^k_{r,p_j} \mid j \leq d, k \leq l\} \times \mathsf{Ord}^{\mathcal{D}'} \times \{\mathsf{cdn}\}
\end{aligned}$$

and $S_+ \subseteq S^\circ_+ \subseteq \mathcal{S}_+ \times \mathcal{S}'_+$ and $R \subseteq S^\circ_+ \times S'_+$

$$\begin{aligned}
\gamma S_+^- \gamma' &\iff \mathsf{n} = \mathsf{n}' \text{ and } j = j', \\
\gamma^\circ R \gamma' &\iff \mathsf{n}^\circ = \mathsf{n}', \text{ and } j^\circ = j', \\
\gamma S_+ \gamma' &\iff \mathsf{n} = \mathsf{n}', \ j = j' \text{ and } \mathsf{ctr}(r) = \mathsf{ctr}'(r')
\end{aligned}$$

for every $\gamma = ((\mathsf{n}, p_j), \mathsf{ctr}, \mathsf{cdn}) \in \mathcal{S}_+$, $\gamma^\circ = ((\mathsf{n}^\circ, x_{p_{j^\circ}}), \mathsf{ctr}^\circ, \mathsf{cdn})$ and $\gamma' = ((\mathsf{n}', \psi^k_{r,p_{j'}}), \mathsf{ctr}', \mathsf{cdn}) \in \mathcal{S}'_+$

Using these relations we get the induction goal (4.9) from Proposition 2.2.2 instantiating $\mathcal{G} = \mathcal{G}(\mathcal{A})$ and $\mathcal{G}' = \mathcal{G}(\psi_{r+1,q})$. The proposition assumes two premises:

$$\mathcal{G}(\mathcal{A}), ((\mathsf{m}, q), \mathsf{ctr}_I, \mathsf{cdn}) \bowtie_{S_+ \sqcup S} \mathcal{G}(\psi_{r+1,q}), ((\mathsf{m}, \psi_{r+1,q}), \mathsf{ctr}'_I, \mathsf{cdn}) \tag{4.14}$$

and:

$$\mathcal{G}(\mathcal{A}), \gamma \bowtie_S \mathcal{G}(\psi_{r+1,q}), \gamma' \tag{4.15}$$

for all $(\gamma, \gamma') \in S_+$. The first premise (4.14) follows from:

$$\mathcal{G}(\mathcal{A}), ((\mathsf{m}, q), \mathsf{ctr}_I, \mathsf{cdn}) \bowtie_{S_+^- \sqcup S} \mathcal{G}(\psi_{r+1,q}), ((\mathsf{m}, \psi_{r+1,q}), \mathsf{ctr}'_I, \mathsf{cdn}) \tag{4.16}$$

in a similar way as (4.12) follows from (4.13) but the argument is even simpler: none of the two above partial games decrements counter for r and r', respectively, before it stops. This is because all the positions with rank r and r' belong to $\mathcal{S}_+$ and $\mathcal{S}'_+$, respectively. Equivalence (4.16) follows from composing:

$$\begin{aligned}\mathcal{G}(\mathcal{A}), ((\mathsf{m}, q), \mathsf{ctr}_I, \mathsf{cdn}) &\bowtie_{S_{r\leq}} \mathcal{G}(\psi_{r,q}), ((\mathsf{m}, \psi_{r,q}), \mathsf{ctr}'_I, \mathsf{cdn}) \\ &\bowtie_{R \sqcup \mathsf{Id}(\mathcal{S}')} \mathcal{G}(\psi_{r+1,q}), ((\mathsf{m}, \psi_{r+1,q}), \mathsf{ctr}'_I, \mathsf{cdn}).\end{aligned}$$

The first line is the induction hypothesis (4.8) applied to $\psi_{r,q}$. The second one follows from the definition of $\psi_{r+1,q}$ which implies that the partial games $\mathcal{G}(\psi_{r,q})|\mathcal{S}^\circ_+ \sqcup \mathcal{S}', (\mathsf{m}, \psi_{r,q})$ and $\mathcal{G}(\psi_{r+1,q})|\mathcal{S}'_+ \sqcup \mathcal{S}', (\mathsf{m}, \psi_{r+1,q})$ are identical until they move to some (n, x_{p_i}) and $(\mathsf{n}, \psi^k_{r+1,p_i})$ (ψ^k_{r+1,p_i} replaces x_{p_i}, where the irrelevant index k depends on the particular occurrence of x_{p_i}). If this happens, the first game stops, whereas the second one runs a counter update (which is trivial since x_{p_i} has standard rank) and moves deterministically to $(\mathsf{n}, \psi^k_{r+1,p_i})$ where it stops. We have:

$$S_{r\leq} \circ (R \sqcup \mathsf{Id}(\mathcal{S}')) = S_+ \sqcup S$$

as in the previous case which allows us to compose the two lines into (4.16).

Towards the second premise (4.15), take $(\gamma, \gamma') \in S_+$ with $\gamma = ((\mathsf{n}, p_j), \mathsf{ctr}, \mathsf{cdn})$ and $\gamma' = ((\mathsf{n}, \psi^k_{r,p_j}), \mathsf{ctr}', \mathsf{cdn})$ such that $\mathsf{ctr}(r) = \mathsf{ctr}'(r')$. We compose:

$$\begin{aligned}\mathcal{G}(\mathcal{A}), ((\mathsf{n}, p_j), \mathsf{ctr}, \mathsf{cdn}) &\bowtie_S \mathcal{G}(\psi_{r+1,p_j}), ((\mathsf{n}, \psi_{r,p_j}), \mathsf{ctr}', \mathsf{cdn}) \\ &\bowtie_{\mathsf{Id}(\mathcal{S}')} \mathcal{G}(\psi_{r+1,q}), ((\mathsf{n}, \psi^k_{r,p_j}), \mathsf{ctr}', \mathsf{cdn}).\end{aligned}$$

The first line is (4.11). The second line follows from our previous analysis: it is not possible to leave $\mathsf{SubFor}(\theta_k)$ from ψ^k_{r,p_j} and thus the reachable part of the game on the right is isomorphic with the game on the left. This completes the proof of Theorem 4.4.5. $\square$

4.5 Facts about the Logic

To demonstrate usefulness of the correspondence between formulae and automata, but also for technical use in further proofs, we shall now establish some facts about the logic (and automata).

4.5.1 Guarded Formulae and Automata

We start showing that without loss of generality formulae are guarded.

Definition 4.5.1. We say that an automaton $\mathcal{A}$ is *guarded* if it does not contain a loop without modal transitions (i.e. a sequence of states $q_1...q_l$ with $q_{i+1} \in \delta(q_i)$ for all $i < l$ and $q_1 \in \delta(q_l)$). A formula φ is guarded if it is guarded when seen as an automaton $\mathcal{A}_\varphi$.

It is relatively easy to see that an automaton can be turned into an equivalent guarded one.

Proposition 4.5.2. *Every countdown automaton can be transformed into an equivalent guarded one.*

Proof. Recall that a loop in a countdown play is called bad for player P if it leads from position v to itself, P owns the rank of r and no position in the loop has a rank greater than r. The maximal length of a sequence of consecutive ε-moves not containing a bad loop in the semantic game $\mathcal{G}(\mathcal{A})$ is bounded by some $k < \omega$ depending on the automaton. Take a new automaton $\mathcal{A}'$ that simulates the original $\mathcal{A}$ and additionally stores in its memory the sequence of all states visited since the last modal move. In case $\mathcal{A}'$ detects a loop bad for P, it stops and P looses immediately. Since the bound k depends only on $\mathcal{A}$ and not on the model, $\mathcal{A}'$ is well defined as it only needs to remember a sequence of states of length at most k. By Proposition 4.3.5, without loosing generality we may require players to use strategies that avoid loops bad for them, so $\mathcal{A}'$ accepts the same language as $\mathcal{A}$. Since $\mathcal{A}'$ is guarded by design, this completes the proof. $\square$

Let us have a look at how the maximal number k of ε-transitions without a bad loop depends on $\mathcal{A}$. In a sequence of ε-transitions without a bad loop, no state q with the most important rank r can repeat. Hence, there could be at most $|\mathsf{rank}^{-1}(r)| \leq |Q|$ visits to r in the sequence. The places where the sequence passes through r decompose it into subsequences that visit fewer ranks. Using these observations one can show by induction on d that if w is such a sequence visiting d ranks then $|w| \leq (|Q|+1)^d$. Since the total number $|\mathcal{R}|$ of ranks of $\mathcal{A}$ is fixed, we get the upper bound:

$$k \leq (|Q|+1)^{|\mathcal{R}|}$$

as $d \leq |\mathcal{R}|$. If the ranks are assigned injectively to states, the above reasoning shows $k \leq 2^{|\mathcal{R}|}$.

On the other hand, we give an exponential lower bound on k, even in the case when the automaton is injectively ranked. For simplicity assume that there is an ε-transition between every two states, that the ranks are given as a finite subset $\{0, ..., n\} \subseteq \omega$ with the usual ordering and that each rank i is assigned to a unique state q_i. Recursively define $w_0 = \varepsilon$ and $w_{i+1} = w_i q_{i+1} w_i$. The sequence w_n has length $2^n - 1$ yet it contains no bad loops and so it witnesses:

$$2^{|\mathcal{R}|} - 1 \leq k.$$

Let us now have a look at a particular case of automata: logical formulae.

Proposition 4.5.3. *Every countdown formula can be transformed into an equivalent guarded one.*

Proof. Every formula can be seen as an automaton, massaged via Proposition 4.5.2 to obtain a guarded one and translated to a new, equivalent formula. Inspecting the translation from Subsection 4.4.2 one could check that it preserves guardedness and so the new formula is guarded, therefore proving Proposition 4.5.3. Instead of doing that, let us give a more direct translation that exploits the game semantics for the logic but avoids abstract automata.

Given a formula φ we modify it inductively (proceeding from leaves to the root of the syntactic tree). Replace every subformula $\theta = \eta^\alpha_k(x_1, ..., x_n).(\psi_1, ..., \psi_n)$ with:

$$\theta' = \eta^\alpha_{k,1}(x_{i,j})_{i,j \leq n}.(\psi_{i,j})_{i,j \leq n}$$

where $\psi_{i,j}$ is obtained from ψ_i by replacing

(i) every guarded x_m with $x_{m,1}$ and

(ii) every other x_m with $\top/\bot$ (respectively) if $j = n$ and $\eta = \nu/\mu$, or with $x_{m,j+1}$ otherwise.

This way, the number of visits to immediate subformulae of the η^α operator without any modal move is counted in the index j. If the index j reaches n and the play moves to some variable $x_{m,n}$, then by the pigeonhole principle some of the formulae $\psi_{i,j}$ visited without a modal move nor a visit to a superformula of θ' must have the same index i. This in turn corresponds to a repeated visit to some ψ_i in the semantic game for the original formula, without modal moves nor visits to a superformula of θ, which implies existence of a loop bad for $\exists$ve/$\forall$dam when $\eta = \nu/\mu$, respectively. $\square$

Note that in the case of scalar formulae, the above construction yields a guarded formula that is also scalar and linear in the size of the original one: we just replace every unguarded occurrence of a variable bound by μ^α (or ν^α) by $\bot$ (or $\top$, respectively).

4.5.2 Pre-modal Phase

The Example 4.3.3 shows that countdown games heavily depend on the counters: the players may need to look at the counter values not only to update them but also to choose moves in the arena. Proposition 4.3.4 gives a very restricted form of counter-independence. Although the guaranteed property of strategies is slightly abstract, in the context of games induced by automata (and in particular, formulae) it can be turned into a more human-friendly form.

For an automaton $\mathcal{A}$ with states Q, a valuation val and a point m_I in a model $\mathcal{M}$, the *pre-modal phase* of the game $\mathcal{G}^{\mathsf{val}}(\mathcal{A}), (\mathsf{m}_I, q_I)$ consists of all pre-modal plays, i.e. plays with no modal move. All the positions accessible in that phase are of the form (m_I, q) for $q \in Q$ and if $\mathcal{A}$ is guarded, then no pre-modal play is longer than $|Q|$. Hence, it follows from Proposition 4.3.4 that:

Proposition 4.5.4. *In every game $\mathcal{G}^{\mathsf{val}}(\mathcal{A}), (\mathsf{m}_I, q_I)$ for a guarded automaton $\mathcal{A}$, the winning player has a pre-modally counter-independent (i.e. counter-independent in the pre-modal phase) winning strategy.*

Since all the positions appearing in the pre-modal phase only have the initial point on the first coordinate, we can identify pre-modal plays π and π' starting in (m, q) and (m', q) for different $\mathsf{m} \neq \mathsf{m}'$ if π equals π' after swapping m and m'. Likewise, we simplify the pre-modal guide:

$$\sigma^I : (\{\mathsf{m}\} \times Q)^{<|Q|} \to \{\mathsf{m}\} \times Q$$

guiding pre-modal σ-plays to:

$$\sigma^I : Q^{<|Q|} \to Q$$

by skipping the redundant first coordinate. Note that the assumption that σ^I guides some winning strategy from point m may imply some conditions on the atomic propositions satisfied by that point.

The next two propositions are a bit technical but will be useful in Sections 4.6 and 4.8. We first prove that strategies with equal pre-modal guides can be synchronized so that they agree on all the pre-modal moves (and not only the positional ones).

Proposition 4.5.5. *Consider two points $\mathsf{m}_0, \mathsf{m}_1$ in a model $\mathcal{M}$, a valuation val and a guarded automaton $\mathcal{A}$. Assume that $\exists$ve wins the game $\mathcal{G}^{\mathsf{val}}(\mathcal{A})$ from m_0 and m_1 with pre-modally counter-independent strategies σ_0 and σ_1, respectively, both guided by the same pre-modal guide σ^I. Then there are winning strategies σ_0', σ_1' guided by σ^I such that:*

- *σ_0' and σ_1' behave the same in the pre-modal phase, up to swapping the points m_0 and m_1, and*

- *for every* $(\mathsf{m}_i, \mathsf{ctr})$ *reachable by a* σ_i'*-play, there are* $(\mathsf{m}_0, \mathsf{ctr}_0)$ *and* $(\mathsf{m}_1, \mathsf{ctr}_1)$ *such that each* $(\mathsf{m}_j, \mathsf{ctr}_j)$ *is reachable by a* σ_j*-play and* $\mathsf{ctr}_i \preccurlyeq_{\mathsf{better}} \mathsf{ctr}$.

A symmetric statement works for $\forall$*dam.*

Proof. By symmetry, we only focus on the case for $\exists$ve. Starting in m_0 or m_1, she can maintain the invariant that for the play π so far, there are π_0 and π_1 consistent with σ_0 and σ_1 respectively, such that (i) all the three plays are (point-wise) equal on her choices of positions and on all $\forall$dam's choices, and (ii) her choices of counter values in π are the maximum of the corresponding choices from π_0 and π_1. This way either she wins in the pre-modal phase or the play reaches a modal move with counter values at least as good for her as after some σ_0- and σ_1-plays, respectively. She may then continue from m_i with the winning strategy σ_i. $\square$

Under some monotonicity conditions, strategies with the same pre-modal guide can be combined into a third one.

Proposition 4.5.6. *Consider three points* $\mathsf{m}_1, \mathsf{m}_2, \mathsf{m}_3$ *in a model* $\mathcal{M}$ *s.t. for every* $\mathsf{a} \in \mathsf{Act}$*, the sets* $S_1^{\mathsf{a}}, S_2^{\mathsf{a}}, S_3^{\mathsf{a}}$ *of their* a*-successors are monotone, i.e.* $S_1^{\mathsf{a}} \subseteq S_2^{\mathsf{a}} \subseteq S_3^{\mathsf{a}}$*; a valuation* val *that does not distinguish* m_i *(i.e.* $\mathsf{m}_i \in \mathsf{val}(x) \iff \mathsf{m}_j \in \mathsf{val}(x)$ *for all* $x \in \mathsf{Var}$*); and a guarded automaton* $\mathcal{A}$*. If a player* P *wins the semantic game* $\mathcal{G}^{\mathsf{val}}(\mathcal{A})$ *from* m_1 *and* m_3 *using strategies* σ_1, σ_3 *guided by the same pre-modal guide* σ^I*, then* P *also wins from* m_2 *with a strategy* σ_2 *guided by* σ^I*.*

Proof. By Proposition 4.5.5, we may assume that σ_1 behaves the same as σ_3 in the pre-modal phase. Initially P may apply the same strategy from m_2, as the point in the model does not matter, or does not change, in the pre-modal phase. Consider any play consistent with this strategy. If P does not win already in the pre-modal phase, the play reaches a modal move, i.e. a configuration (m_2, q) with $q \in Q$ such that $\delta(q) = (\mathsf{a}, p)$. If the state q is owned by P then P may continue with σ_1, and if q is owned by P's opponent then P may continue with σ_3. $\square$

4.6 Vectorial vs. Scalar Calculus

In this section we investigate the relation between scalar and vectorial formulae. We have already seen with Example 4.2.5 that unlike with standard fixpoints, the Bekić principle is not valid in the countdown setting. Interestingly, scalar formulae correspond to automata with a simple syntactic restriction.

Proposition 4.6.1. *Scalar countdown formulae and automata where every two states have different ranks have equal expressive power.*

Proof. Inspecting the translations between formulae and automata from Subsections 4.4.1 and 4.4.2, it is evident that injectively ranked automata are translated to scalar formulae, and that, although in our translation the choice of the ranks function is not deterministic, every scalar formula can be translated to an injectively ranked automaton. □

Since the Bekić principle fails, a natural question is whether there is another way of transforming vectorial formulae to scalar form (or, equivalently, arbitrary countdown automata to injectively ranked ones). We shall give a negative answer in Theorem 4.6.2. However, before we proceed, let us analyze the following example, which shows that scalar formulae are more expressive than they may seem, covering in particular the property from Example 4.2.5.

4.6.1 Languages of Unbounded Infixes

Fix a regular language of finite words $L \subseteq \Gamma^*$. Let $\mathcal{U}(L) \subseteq \Gamma^\omega$ be the language of all infinite words that contain arbitrarily long infixes from L. For instance, the language from Example 4.2.5 is $\mathcal{U}(\mathsf{a}^*)$. We shall now show that $\mathcal{U}(L)$ can be defined in the countdown μ-calculus, first by a vectorial formula, then by a scalar one.

Consider a finite deterministic automaton $\mathcal{A} = (Q, \delta, q_I, F)$ that recognizes L. Let $\delta^+ : \Gamma^+ \times Q \to Q$ be the unique inductive extension of the transition function $\delta : \Gamma \times Q \to Q$ to nonempty words. Define:

$$K_{p,q} = \{w \in \Gamma^+ \mid \delta^+(w, p) = q\}$$

the (regular) language of nonempty words leading from p to q in $\mathcal{A}$, and let $K_{p,F}$ denote the union $\bigcup_{q\in F} K_{p,q}$. By the pigeonhole principle we have $\mathcal{U}(L) = \bigcup_{q\in Q} \mathcal{U}_q(L)$, where $\mathcal{U}_q(L) \subseteq \Gamma^\omega$ consists of words such that for every $n < \omega$, w has an infix $w_n = v_I u_1 ... u_n v_F \in L$ s.t. (i) $v_I \in K_{q_I,q}$, (ii) $u_1, ..., u_n \in K_{q,q}$, and (iii) $v_F \in K_{q,F}$. Then $\mathcal{U}_q(L)$ can be defined by a vectorial formula:

$$\mathcal{U}_q(L) = [\![\nu_1^\omega(x_1, x_2).(\langle \Gamma^* K_{q_I,q}\rangle x_2, \langle K_{q,q}\rangle x_2 \wedge \langle K_{q,F}\rangle \top)]\!]$$

where $\langle K\rangle\psi$ is the formula as explained in Example 4.2.5. Indeed, the corresponding semantic game on a word w proceeds as follows:

1. $\forall$dam chooses a number $n < \omega$ as the value of his only counter,
2. $\exists$ve skips a prefix $v_0 v_I \in \Gamma^* K_{q_I,q}$ of w,
3. $\forall$dam decrements his counter;
4. $\exists$ve keeps moving through $u_1, u_2, ... \in K_{q,q}$ so that after each step, some state in F is reachable from q by some prefix of the remaining

word. After each such choice of u_i $\forall$dam has to decrement his counter, and so $\exists$ve wins iff she can make at least $n-1$ such steps.

The two different stages in which $\forall$dam's counter is decremented reflect the two-phase dynamics of the game: first $\forall$dam challenges $\exists$ve with a number, and then $\exists$ve shows that she can provide an infix long enough.

It is more tricky to define the language $\mathcal{U}_q(L)$ with a scalar formula, but it turns out to be possible. To this end, observe that without loss of generality we may restrict attention to words w such that:

1. the infixes $w_n \in L$ start arbitrarily far in w;
2. each w_n can be decomposed as $v_I u_1 ... u_n v_F \in L$ s.t. (i) $v_I \in K_{q_I,q}$, (ii) $u_1, ..., u_n \in K_{q,q}$, (iii) $v_F \in K_{q,F}$, and additionally *(iv) all u_i begin with the same letter* $\mathsf{a} \in \Gamma$;
3. there are at least two distinct letters $\mathsf{a}, \mathsf{b} \in \Gamma$ that appear infinitely often in w;
4. the first letter of w is b.

Indeed, for (1) note that otherwise w_n start in the same position k for all n large enough. But then even the stronger property "There exists a position k such that the run of $\mathcal{A}$ from k visits q and F infinitely often" holds, and this is easily definable by a fixpoint formula.

Item (2) follows from the pigeonhole principle and the observation that in $w_{n\times|\Gamma|} = v_I u_1 ... u_{n\times|\Gamma|} v_F$ at least n u_i's begin with the same letter.

For (3) observe that otherwise w has a suffix a^ω for some $\mathsf{a} \in \Gamma$, in which case membership in $\mathcal{U}_q(L)$ is definable by a fixpoint formula. This is because an ultimately periodic word is bisimilar to a finite model, and so every monotone map reaches its fixpoints in finitely many steps, meaning that the countdown operator ν^ω is equivalent to ν^∞.

Finally, for (4) note that the language $\mathcal{U}_q(L)$ is closed under adding and removing finite prefixes, and so if a formula φ defines $\mathcal{U}_q(L) \cap \mathsf{b}\Gamma^\omega$, then the formula $\langle\Gamma^*\rangle(\langle\mathsf{b}\rangle\top \wedge \varphi)$ defines $\mathcal{U}_q(L)$.

With this in mind, define:

$$\varphi = \nu^\omega x.(\langle\mathsf{b}\rangle\top \wedge \langle\Gamma^* K_{q_I,q}\rangle(\langle\mathsf{a}\rangle\top \wedge x)) \vee (\langle K_{q,q}\rangle(\langle\mathsf{a}\rangle\top \wedge x) \wedge \langle\mathsf{a}\rangle\top \wedge \langle K_{q,F}\rangle\top).$$

Note how $\langle\mathsf{b}\rangle\top \wedge x$ and $\langle\mathsf{a}\rangle\top \wedge x$ replace x_1 and x_2 from the vectorial formula. Consider the corresponding semantic game on a word w. Consider configurations of the game with the main disjunction as the formula component. Every infinite play of the game must visit such configurations infinitely often. In such a configuration, if the next letter in the model is either a or b then $\exists$ve must choose the right or left disjunct, respectively. In particular, once the game reaches a configuration where $\langle\mathsf{a}\rangle\top$ holds, it must also hold every

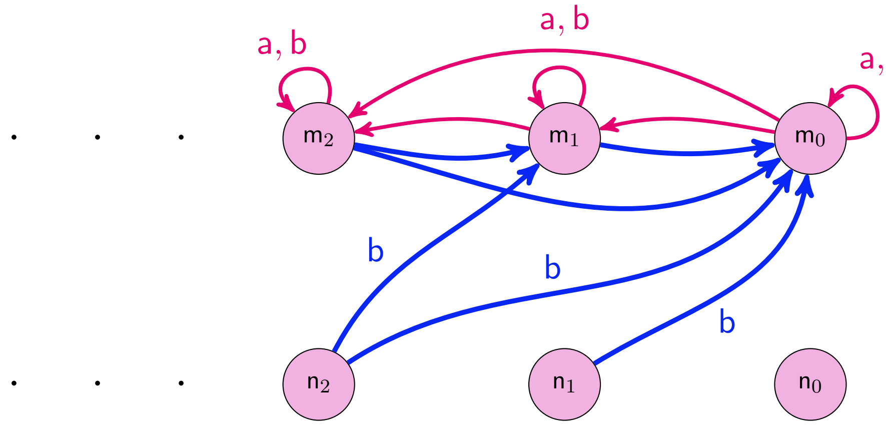

Figure 4.1: The model $\mathcal{M}$. Blue arrows represent edges labeled both with a and b, and pink arrows are edges labeled only with b.

time the variable x is unraveled in the future. As a result, ∃ve wins from a configuration where $\langle \mathsf{a} \rangle \top$ holds against ∀dam's counter $n < \omega$ iff there is $u_1 ... u_{n+1} v_F$ starting in the current position such that $u_1, ..., u_{n+1} \in K_{q,q}$, each u_i starts with a, and $v_F \in K_{q,F}$. Moreover, ∃ve wins from a position where $\langle \mathsf{b} \rangle \top$ holds, against ∀dam's $n+1 < \omega$, iff there is $v_I \in \Gamma^* K_{q_I,q}$ starting in the current position such that the next position after v_I satisfies $\langle \mathsf{a} \rangle \top$ and ∃ve wins from there against n. Putting this together, we get that ∃ve wins from a position satisfying $\langle \mathsf{b} \rangle \top$ against n iff there is $v_I u_1 ... u_n v_F = w_n$ as in condition (2) above. Since the game starts with ∀dam choosing an arbitrary $n < \omega$, it follows that indeed φ defines $\mathcal{U}_q(L)$.

4.6.2 Greater Expressive Power of the Vectorial Calculus

We now show an example of a property that is definable in the vectorial countdown calculus but not in the scalar one.

Fixing $\mathsf{Act} = \{\mathsf{a}, \mathsf{b}\}$, consider a model $\mathcal{M} = (M, \xrightarrow{\mathsf{a}}, \xrightarrow{\mathsf{b}})$ with points $M = \{\mathsf{m}_i, \mathsf{n}_i \mid i < \omega\}$, and with exactly the edges: $\mathsf{m}_i \xrightarrow{\mathsf{a}} \mathsf{m}_j$, $\mathsf{n}_i \xrightarrow{\mathsf{a}} \mathsf{m}_j$ and $\mathsf{n}_i \xrightarrow{\mathsf{b}} \mathsf{m}_j$ for all $i > j$; and $\mathsf{m}_i \xrightarrow{\mathsf{b}} \mathsf{m}_j$ for all i and j. Note that the relation $\xrightarrow{\mathsf{a}}$ is a subset of $\xrightarrow{\mathsf{b}}$. The model is shown in Fig. 4.1.

Consider the vectorial sentence $\nu_1^{\omega}(x_1, x_2).(\langle \mathsf{b} \rangle x_2, \langle \mathsf{a} \rangle x_2)$. This describes the property *there are arbitrarily long paths with labels in* ba^*, and so it is true in all points m_i and false in all points n_i. The following result immediately implies that this property cannot be defined in the scalar countdown calculus:

Theorem 4.6.2. *For every scalar sentence φ, there exists $i < \omega$ such that:*

$$\mathsf{m}_i \in [\![\varphi]\!] \iff \mathsf{n}_i \in [\![\varphi]\!].$$

Proof. Observe that since scalar sentences are closed under negation, it is enough to prove that for every scalar φ there is a $\kappa_\varphi < \omega$ such that for all $i > \kappa_\varphi$:

$$\mathsf{m}_i \in [\![\varphi]\!] \implies \mathsf{n}_i \in [\![\varphi]\!] \tag{$\star$}$$

Moreover, since every scalar formula can be transformed into an equivalent *guarded* formula that is also scalar, it suffices to prove ($\star$) for guarded formulae. For the rest of the proof, we fix a scalar sentence φ and denote $\mathcal{G} = \mathcal{G}(\varphi) = (V, E, \mathsf{rank}, \mathsf{ctr}_I)$. Thanks to Proposition 4.6.1 we assume φ when seen as an automaton is injectively ranked and Proposition 4.5.3 allows us to assume that φ is guarded. Let us start with an easy fact.

Proposition 4.6.3. *There exists some $N < \omega$ such that for all $N \leq i < j$:*

$$\mathsf{m}_i \in [\![\varphi]\!] \iff \mathsf{m}_j \in [\![\varphi]\!]$$

and if $\exists$ve wins the corresponding evaluation games then she does so with pre-modally counter-independent strategies σ_{m_i} and σ_{m_j} with the same pre-modal guide σ^I that does not depend on i, j.

Proof. Note that the relations $\xrightarrow{\mathsf{a}}$ and $\xrightarrow{\mathsf{b}}$ are monotone, i.e. the bigger i, the more a- and b-successors m_i has. On the other hand, there are only finitely many possible pre-modal guides, so by the pigeonhole principle if $\exists$ve wins from m_i for arbitrarily big i, some pre-modal guide σ^I is used for arbitrarily big i. Thus, by Proposition 4.5.6 she can use σ^I to win for *all i big enough*. □

Towards ($\star$), assume that $\mathsf{m}_i \in [\![\varphi]\!]$ for all i big enough (otherwise, by Proposition 4.6.3, φ is *false* in m_i for all i big enough, which trivially implies ($\star$)) and denote by N the least number for which Proposition 4.6.3 holds. We show that without losing generality the strategies have some nice properties:

Proposition 4.6.4. *The family $\{\sigma_{\mathsf{m}_i} \mid i < \omega\}$ of strategies can be massaged so that for every $i > N$, if a σ_{m_i}-play visits a modal position for the first time and it has the shape $(\mathsf{m}_i, \langle \mathsf{a} \rangle \psi)$, then σ_{m_i} chooses a point m_j for some $j < N$.*

Proof. By Proposition 4.6.3, σ_{m_i} and σ_{m_N} have the same pre-modal guide σ^I. Therefore, by Proposition 4.5.5, there is a strategy σ'_{m_i} winning from m_i guided by the same σ^I and only reaching pre-modal configurations at least as good for $\exists$ve as the ones reachable by σ_{m_N}. Then, whenever $(\mathsf{m}_i, \langle \mathsf{a} \rangle \psi)$ is reached in a pre-modal σ'_{m_i}-play, by the monotonicity of $\xrightarrow{\mathsf{a}}$ and the assumption that $N < i$, $\exists$ve may just continue with σ_{m_N}. Moreover, since σ_{m_N} is a legitimate strategy, it must pick a point m_j for some $j < N$, as desired. □

Denote by $\mathcal{B}$ the phase of the game $\mathcal{G}(\varphi)$ that consists of plays of the shape $\pi = \xi\rho$ such that the play ξ is pre-modal and leads to $\langle \mathsf{b} \rangle$ and ρ (starting with an $\exists$ve's choice of a $\xrightarrow{\mathsf{b}}$-successor of the initial point m_i) does not visit (i) a formula beginning with $[\mathsf{a}], [\mathsf{b}]$ or $\langle \mathsf{b} \rangle$, nor (ii) a formula with a rank that was visited in the pre-modal phase (that is, in ξ) with a possible exception for its last configuration:

$$\mathcal{B} = \left\{ \xi\rho \;\middle|\; \begin{array}{l} \xi \text{ is pre-modal and leads to } \langle \mathsf{b} \rangle;\ \rho \text{ is nonempty and has no} \\ \text{proper prefix visiting } [\mathsf{a}], [\mathsf{b}], \langle \mathsf{b} \rangle \text{ nor a rank visited in } \xi \end{array} \right\}$$

where by a visit to a modal operator $\heartsuit$ we mean a visit to a formula $\heartsuit\psi$ beginning with it. Note that the only modal moves that start and end in $\mathcal{B}$ are those corresponding to $\langle \mathsf{a} \rangle$. The next step is the following further enhancement of the strategies:

Proposition 4.6.5. *The family $\{\sigma_{\mathsf{m}_i} \mid i < \omega\}$ of strategies can be massaged so that there exists a finite bound $k_{\mathsf{max}} < \omega$ such that no σ_{m_i}-play $\pi \in \mathcal{B}$ contains more than k_{max} modal moves.*

Before proving the above proposition, let us demonstrate how it implies $(\star)$. Put:

$$\kappa_\varphi = k_{\mathsf{max}} + N + 1$$

where k_{max} is the bound from Proposition 4.6.5. We show that $\mathsf{n}_i \in [\![\varphi]\!]$ for every $i > \kappa_\varphi$. To this end, consider the strategy σ_i with $\mathsf{above}_{+1}(\sigma_{\mathsf{m}_i}, \sigma_i)$, i.e. the strategy 1-above σ_{m_i} which exists by Proposition 2.2.4. In the pre-modal phase of the evaluation game from (n_i, φ), use σ_i (recall that we identify pre-modal plays starting in different (m_i, φ) and (n_i, φ) if they are equal up to swapping positions (m_i, θ) and (n_i, θ) for all θ). Since m_i and n_i have the same $\xrightarrow{\mathsf{a}}$-successors, if a play visits $\langle \mathsf{a} \rangle$ or $[\mathsf{a}]$, $\exists$ve may continue with σ_i and win. The same is true for $[\mathsf{b}]$, as every $\xrightarrow{\mathsf{b}}$-successor of n_i is also a $\xrightarrow{\mathsf{b}}$-successor of m_i.

The only interesting case is when a play reaches $\langle \mathsf{b} \rangle$ and σ_i chooses $\mathsf{m}_{j'}$ which cannot be chosen from i, i.e. $j' \geq i$ (if $j' < i$ then $\mathsf{m}_{j'}$ is a $\xrightarrow{\mathsf{b}}$-successor of n_i, so $\exists$ve may use σ_i). In this case, $\exists$ve may choose m_j where $j = k_{\mathsf{max}} + N$ and play maintaining the invariant that for the current play π, as long as it belongs to $\mathcal{B}$, there is a σ_i-play π' in $\mathcal{B}$ such that:

1. all subformulae and ordinals are the same in π and π',
2. for the last points m_j and $\mathsf{m}_{j'}$ of π and π', respectively, we have:

   $$k + N \leq j \leq j' \qquad \text{and} \qquad j < i$$

   where k is the bound on the number of modal moves that can be made in a σ_i-play extending π' without leaving $\mathcal{B}$.

It is straightforward to maintain the invariant on ε-transitions and when counter values are updated. If π ends with a visit to $\langle \mathsf{a} \rangle$ (so $\exists$ve has to pick an $\xrightarrow{\mathsf{a}}$-successor of m_j) and $\sigma_i(\pi')$ dictates the choice of $\mathsf{m}_{j''}$ for some $j'' < \omega$ then either:

- $j > j''$, hence $\mathsf{m}_j \xrightarrow{\mathsf{a}} \mathsf{m}_{j''}$ and $\exists$ve wins using $\pi\rho \mapsto \sigma_i(\pi'\rho)$, or
- $j \leq j''$, which combined with item (2) of the invariant gives:

$$k + N \leq j \leq j''.$$

  Since π' extended with the choice of $\mathsf{m}_{j''}$ ends with a modal move, we are left with at most $k-1$ possible modal moves in $\mathcal{B}$, so the choice of m_{j-1} preserves the invariant, as $k-1+N \leq j-1 \leq j''$ and $j-1 < j < i$.

Since φ is guarded, the maximal number of consecutive ε-transitions in a play is bounded by $|\mathsf{SubFor}(\varphi)|$ (which corresponds to $2 \cdot |\mathsf{SubFor}(\varphi)|$ moves, including the countdown and the positional ones). Since the number of modal moves before the end of $\mathcal{B}$ is bounded by k_{max}, after at most:

$$2 \cdot k_{\mathsf{max}} \cdot |\mathsf{SubFor}(\varphi)| < \omega$$

moves either the game ends in $\mathcal{B}$ or eventually visits (i) $[\mathsf{a}], [\mathsf{b}]$ or $\langle \mathsf{b} \rangle$ or (ii) a rank seen in the pre-modal phase. If the game ends in $\mathcal{B}$ then thanks to item (1) of the invariant $\exists$ve wins (as the strategy σ_i is winning).

In the other case with (i) or (ii), the invariant implies that for the current play π there is a σ_i-play π' satisfying (1) and (2).

When (i) the play visits $[\mathsf{a}], [\mathsf{b}]$ or $\langle \mathsf{b} \rangle$, $\exists$ve may continue as with σ_i, i.e. using σ_i' defined as:

$$\sigma_i'(\pi\rho) = \sigma_i(\pi'\rho)$$

for all ρ. If π leads to $[\mathsf{a}]$ then σ_i' is legal due to monotonicity of $\xrightarrow{\mathsf{a}}$ and inequality $j \leq j'$ guaranteed by item (2) of the invariant. Similarly, since m_j and $\mathsf{m}_{j'}$ have the same $\xrightarrow{\mathsf{b}}$-successors, σ_i' is legal also if the play leads to $\langle \mathsf{b} \rangle$ or $[\mathsf{b}]$. Since σ_i'-plays differ from σ_i-plays only by a finite prefix π, $\exists$ve wins.

The remaining case (ii) is when the game moves to a subformula ψ such that $\mathsf{rank}(\psi) = r$ was visited in the pre-modal phase. Since φ is scalar (and so by Proposition 4.6.1 injectively ranked when seen as an automaton), this implies that actually the same *formula* ψ must have been visited in the pre-modal phase.

Denote by π_ψ and π'_ψ the pre-modal prefixes of π and π', respectively, ending just after the counter update corresponding to the first visit in ψ (in particular, π_ψ and π'_ψ end with a positional configuration). By item (1) of the invariant, the plays π and π' (π_ψ and π'_ψ) lead to the same counter assignment ctr (ctr_ψ, respectively). Since π'_ψ is a σ_i-play leading to $((\mathsf{m}_i, \psi), \mathsf{ctr}_\psi, \mathsf{psn})$ and by definition σ_i is a strategy 1-above σ_{m_i}, there exists a σ_{m_i}-play π_ψ^-

leading to $((\mathsf{m}_i, \psi), \mathsf{ctr}^-_\psi, \mathsf{psn})$ with $\mathsf{above}_{+1}(\mathsf{ctr}^-_\psi, \mathsf{ctr}_\psi)$. Clearly, $\exists$ve can win from $((\mathsf{m}_i, \psi), \mathsf{ctr}^-_\psi, \mathsf{psn})$ by just continuing the σ_{m_i}-play, i.e. using σ defined as:

$$\sigma(\rho) = \sigma_{\mathsf{m}_i}(\pi^-_\psi \rho)$$

for all ρ. We claim that also:

$$\exists\text{ve wins from } ((\mathsf{m}_j, \psi), \mathsf{ctr}^-_\psi, \mathsf{psn}) \tag{4.17}$$

for every $N \leq j < i$. Indeed, she can use the same σ:

1. In the pre-modal phase σ is legal, as both games start with the same subformula ψ.

2. We justify that σ is valid also when after a pre-modal play ρ the game reaches a formula θ that begins with a modal operator. Observe that since ρ is pre-modal it does not change the point, meaning that it leads from (m_j, ψ) to (m_j, θ) and from (m_i, ψ) to (m_i, θ). Consider cases:

   - If θ begins with $\langle \mathsf{b} \rangle$ or $[\mathsf{b}]$ then the possible moves from the position (m_j, θ) are the same as from (m_i, θ). This is because the points m_i and m_j have the same $\xrightarrow{\mathsf{b}}$-successors. Therefore, $\exists$ve can continue from (m_j, θ) as if she started from (m_i, θ).
   - Similarly, $j \leq i$ implies that every $\xrightarrow{\mathsf{a}}$-successor of m_j is an $\xrightarrow{\mathsf{a}}$-successor of m_i. Hence, if θ begins with $[\mathsf{a}]$ then σ_{m_i} can be used to win against every $\forall$dam's choice of an $\xrightarrow{\mathsf{a}}$-successor of m_j.
   - The remaining case is when θ begins with $\langle \mathsf{a} \rangle$. By Proposition 4.6.4, if in the first modal step of a σ_{m_i}-play $\exists$ve has to provide an $\xrightarrow{\mathsf{a}}$-successor m_k of m_i, then σ_{m_i} chooses some m_k with $k < N$. By definition of σ (as π^-_ψ is pre-modal) the same is true for σ. It follows from $N \leq j$ that the choice given by σ is legal from m_j.

This proves (4.17).

Note that since by definition the end of π (and π') is the first time when the game revisits some rank seen in the pre-modal phase, we have:

$$\mathsf{ctr}_\psi(r') = \mathsf{ctr}(r') \tag{4.18}$$

for every nonstandard rank $r' \geq r$ where r is the rank of ψ. If there was $r' \geq r$ with $\mathsf{ctr}_\psi(r') \neq \mathsf{ctr}(r')$ this would mean that the the counter for r' was either decremented or reset between π_ψ and π. A reset is only possible upon a visit to an even more important $r'' > r'$ and a decrement upon a visit to r'.

Consider two cases. First, if $r = r'$ and the corresponding counter gets decremented then there must be a visit to ψ between π_ψ and π. This is because φ is scalar and thus injectively ranked, so the only possibility to visit r is by visiting ψ. However, since φ is guarded it is not possible to visit ψ twice in the pre-modal phase. And since the end of π is the first time when r is visited in the post-modal phase, the post-modal part of the play between π_ψ and π does not visit r either.

In the remaining case a rank strictly greater than r must have been visited between π_ψ and π. Since no subformula of ψ has a rank greater than r this requires a visit to some formula that is not a subformula of ψ. But the only way to leave $\mathsf{SubFor}(\psi)$ is to unravel a variable bound in some superformula θ of ψ. This, however, could not happen. Since ψ was visited in the pre-modal phase, so was every its superformula, including θ. The rest of the argument is the same as in the previous case: guardedness of ψ implies that θ could not have been revisited in the pre-modal phase. And if θ was visited in some strict post-modal prefix of π this would contradict the assumption that the end of π is the first time when a rank seen in the pre-modal phase is revisited: $\mathsf{rank}(\theta)$ would be revisited earlier. This completes the proof of (4.18).

To finish the proof of ($\star$) we need to show how $\exists$ve can win once π has been played and the game reached a countdown configuration $((\mathsf{m}_j, \psi), \mathsf{ctr}, \mathsf{cdn})$. By item (2) of the invariant, $N \leq j < i$. We show how $\exists$ve can guarantee that for every result ctr' of the upcoming counter update from ctr:

$$\mathsf{ctr}_\psi^- \preccurlyeq_{\mathsf{better}} \mathsf{ctr}' \tag{4.19}$$

which implies:

$$((\mathsf{m}_j, \psi), \mathsf{ctr}_\psi^-, \mathsf{psn}) \preccurlyeq_{\mathsf{better}} ((\mathsf{m}_j, \psi), \mathsf{ctr}', \mathsf{psn}).$$

so that she can win after the update thanks to (4.17).

Note that in the update from ctr to ctr' counters for all the ranks $r' < r$ are reset to their initial values (as they were reset when moving to ctr_ψ) and counters for all $r' > r$ are not changed (and by (4.18) their values will be equal to the ones in ctr_ψ). That is:

$$\mathsf{ctr}'(r') = \mathsf{ctr}_\psi(r')$$

for all $r' \neq r$, i.e. the only possible difference between ctr' and ctr_ψ is in the value for r. On the other hand recall that by definition $\mathsf{above}_{+1}(\mathsf{ctr}_\psi^-, \mathsf{ctr}_\psi)$ and so in particular:

$$\mathsf{ctr}_\psi^- \preccurlyeq_{\mathsf{better}} \mathsf{ctr}_\psi.$$

Consider cases:

- If r is standard it has no corresponding counter and so the update deterministically leads to ctr' with:

$$\mathsf{ctr}_\psi^- \preccurlyeq_{\mathsf{better}} \mathsf{ctr}_\psi = \mathsf{ctr}'$$

  which implies (4.19).

- If r is nonstandard and belongs to $\forall$dam, he has to pick $\mathsf{ctr}'(r) < \mathsf{ctr}(r)$ and all the other counters in ctr' will be the same as in ctr_ψ. Since by (4.18) $\mathsf{ctr}(r) = \mathsf{ctr}_\psi(r)$ it follows that:

$$\mathsf{ctr}_\psi^- \preccurlyeq_{\mathsf{better}} \mathsf{ctr}_\psi \preccurlyeq_{\mathsf{better}} \mathsf{ctr}'$$

  (and hence (4.19)) for every possible ctr'.

- If r is nonstandard and belongs to $\exists$ve, we need to provide her with a valid choice of ctr'. Define:

$$\mathsf{ctr}'(r') = \begin{cases} \mathsf{ctr}_\psi(r') & \text{if } r' > r, \\ \mathsf{ctr}_\psi^-(r) & \text{if } r' = r, \\ \mathsf{ctr}_\psi(r') & \text{if } r' < r. \end{cases}$$

  Such ctr' is a legal update from ctr, because:

$$\mathsf{ctr}(r) = \mathsf{ctr}_\psi(r) = \mathsf{ctr}_\psi^-(r) + 1.$$

  The first equality is (4.18). The second follows from $\mathsf{above}_{+1}(\mathsf{ctr}_\psi^-, \mathsf{ctr}_\psi)$ and the observation that $\mathsf{ctr}_\psi(r) \neq \mathsf{ctr}_I(r)$ because ctr_ψ is the counter assignment immediately after an update corresponding to ψ (which has rank r). It follows from the definition that:

$$\mathsf{ctr}_\psi^- \preccurlyeq_{\mathsf{better}} \mathsf{ctr}'.$$

  That is, we get (4.19).

This completes the proof of ($\star$) from Proposition 4.6.5.

It remains to prove Proposition 4.6.5, i.e. to refine strategies σ_{m_i} to obtain a finite bound $k_{\mathsf{max}} < \omega$ on the number of modal moves in a play in the phase $\mathcal{B}$. We will show a stronger fact: no play $\pi \in \mathcal{B}$ visits the same formula of shape $\langle \mathsf{a} \rangle \psi$ twice. Then, Proposition 4.6.5 follows with the bound $k_{\mathsf{max}} = |\mathsf{SubFor}(\varphi)|$ because a visit to any of the other modal operators ($[\mathsf{a}]$, $\langle \mathsf{b} \rangle$ or $[\mathsf{b}]$) ends $\mathcal{B}$.

Before we go into the somewhat technical details, let us sketch the core idea of the proof which splits into two steps. First, we show that if instead of updating the counters during $\mathcal{B}$ the players postpone the updates and only decrement the counters once, upon leaving $\mathcal{B}$, this does not change the winner of the game. Second, we use this equivalence to massage σ_{m_i} so that instead of performing a sequence:

$$(\mathsf{m}_j, \langle \mathsf{a} \rangle \psi) \to (\mathsf{m}_{j'}, \psi) \to \ \ldots\ \to (\mathsf{m}_l, \langle \mathsf{a} \rangle \psi) \to (\mathsf{m}_{l'}, \psi) \in V^+$$

of modal moves corresponding to $\langle \mathsf{a} \rangle \psi$, $\exists$ve immediately goes to the *last* point $(\mathsf{m}_j, \langle \mathsf{a} \rangle \psi) \to (\mathsf{m}_{l'}, \psi)$. This is possible thanks to transitivity and well-foundedness of $\xrightarrow{\mathsf{a}}$ and avoids repetitions of $\langle \mathsf{a} \rangle \psi$.

To prove the claim, it is enough if for every *minimal* (and therefore necessarily ending with a first modal move, corresponding to $\langle \mathsf{b} \rangle$) $\pi_I \in \mathcal{B}$ we refine σ_{m_i} to a strategy σ_{π_I} so that:

1. σ_{π_I} does not allow to visit any $\langle \mathsf{a} \rangle \psi$ twice in any play $\rho \in \mathcal{B}_{\pi_I}$ and
2. the behavior on all other plays is not changed, meaning that $\sigma_{\pi_I}(\rho) = \sigma_{\mathsf{m}_i}(\rho)$ for every ρ without π_I as a prefix.

If we do that for every minimal $\pi_I \in \mathcal{B}$, we may combine all such refined strategies into one:

$$\sigma_{\mathcal{B}}(\rho) = \begin{cases} \sigma_{\pi_I}(\rho) & \text{if } \pi_I \text{ is the minimal prefix of } \rho \text{ from } \mathcal{B}, \\ \sigma_{\mathsf{m}_i}(\rho) & \rho \text{ has no prefix in } \mathcal{B}; \end{cases}$$

that avoids repetitions of each $\langle \mathsf{a} \rangle \psi$ in *every* $\pi \in \mathcal{B}$.

Towards such a refinement of σ_{m_i}, fix a minimal $\pi_I \in \mathcal{B}$ leading to a winning countdown configuration $\gamma = ((\mathsf{m}_z, \theta_z), \mathsf{ctr}_z, \mathsf{cdn})$. Denote $v_z = (\mathsf{m}_z, \theta_z) \in V$ and:

$$\mathcal{B}^\circ = \{\rho \mid \pi_I \rho \in \mathcal{B}\}.$$

To get our desired σ_{π_I} it suffices to construct a winning strategy for $\mathcal{G}, \gamma$ that avoids repetitions of each $\langle \mathsf{a} \rangle \psi$ in every $\pi \in \mathcal{B}^\circ$ and behaves as σ_{m_i} on plays without π_I as a prefix.

Note that membership in $\mathcal{B}^\circ$ only depends on the underlying positions. Let $V^\circ \subseteq V$ be the set of all the positions of shape (m, ξ) with ξ either (i) beginning with $[\mathsf{a}]$, $[\mathsf{b}]$ or $\langle \mathsf{b} \rangle$ or (ii) having a rank that was visited in a proper prefix of π_I. If $v_z \in V^\circ$, then $\mathcal{B}^\circ = \{\varepsilon\}$ and there is nothing to prove so let us focus on the other case $v_z \notin V^\circ$. Then:

$$\pi \in \mathcal{B}^\circ \iff \text{no proper prefix of } \pi \text{ visits } V^\circ$$

for all plays π in $\mathcal{G}, \gamma$.

We want to consider partial functions $f : V^* \to V$ that are candidates for a $\mathcal{B}^\circ$-guide of some winning strategy σ in $\mathcal{G}, \gamma$. In order to guide σ in $\mathcal{B}^\circ$, f needs to satisfy some basic conditions, e.g. if it allows to traverse vertices $\overline{v} = v_1 ... v_k \in (V - V^\circ)^*$ and v_k belongs to $\exists$ve then f dictates a legal move $v_{k+1} = f(\overline{v})$. To capture such conditions formally, consider the simple parity game $\widetilde{\mathcal{G}} = (V, E, \mathsf{rank})$ obtained from $\mathcal{G}$ by making all the ranks standard. To avoid confusion, we will call simple parity plays $\overline{v} \in V^*$ in $\widetilde{\mathcal{G}}$ *paths* and reserve the term *plays* for $\mathcal{G}$.

Note that since $\widetilde{\mathcal{G}}$ is obtained from $\mathcal{G}$ by forgetting the counters, if a play $\pi \in \mathcal{B}^\circ$ visits positions $\mathsf{pos}(\pi) = \overline{v}$ then $\overline{v}$ is a path in $\widetilde{\mathcal{G}}, v_z$ not visiting V°, with a possible exception for the last position. We also have the opposite. That is, for every $\overline{v} \in V^*$:

$$\overline{v} \text{ is a path in } \widetilde{\mathcal{G}} | V^\circ, v_z \iff \overline{v} \in \mathsf{pos}[\mathcal{B}^\circ]. \tag{4.20}$$

For the missing left to right implication consider a path $\overline{v}$ in $\widetilde{\mathcal{G}}|V^\circ, v_z$. We need to extend $\overline{v}$ to a play in $\mathcal{G}, \gamma$, i.e. provide consistent choices of counter values for both players. It suffices if whenever the owner of rank r has to pick value α, (s)he chooses the number of the remaining visits to r in $\overline{v}$ (so that (s)he is not stuck before traversing $\overline{v}$). Such number is not greater than $|\overline{v}|$ and hence finite. Therefore, to show that the above strategy is legal it is enough to prove that the value before the first update (i.e. in ctr_z) is a limit one. Since $\overline{v} \in (V - V^\circ)^*(V^\circ + \varepsilon)$, no position in $\overline{v}$ other than the last one (for which we do not need to update the counter to prove the claim) has a rank that was visited in π_I. It follows that all the counters to be updated have initial, and hence limit values in ctr_z.

Since every modal move over $\xrightarrow{\mathsf{b}}$ leaves $\mathcal{B}^\circ$, it follows that all the positions visited in $\pi \in \mathcal{B}^\circ$ are of the form (m_k, ψ) for $k \leq z$ (because $\xrightarrow{\mathsf{a}}$ only leads to points with a strictly smaller index and π_I leads to (m_z, θ_z)) and no such position repeats in $\mathcal{B}^\circ$ (by guardedness of φ and acyclicity of $\xrightarrow{\mathsf{a}}$). It follows that the set $\mathsf{pos}[\mathcal{B}^\circ]$ is finite:

$$|\mathsf{pos}[\mathcal{B}^\circ]| < \omega \tag{4.21}$$

and in particular plays in $\mathcal{B}^\circ$ have length bounded by a finite constant. Combining (4.20) and (4.21) we get:

$$\text{The set of paths in } \widetilde{\mathcal{G}}|V^\circ, v_z \text{ is finite.} \tag{4.22}$$

We call a partial function $f : V^* \to V$, thought of as a candidate for a $\mathcal{B}^\circ$-guide of a winning strategy for $\mathcal{G}, \gamma$, a *proto-strategy* if f is a non-loosing strategy in $\widetilde{\mathcal{G}}|V^\circ, v_z$. Every f that is a $\mathcal{B}^\circ$-guide for some winning σ in $\mathcal{G}, \gamma$ is a proto-strategy. Observe that for every f-path $\overline{v} \in (V - V^\circ)^*$ starting from v_z there exists a σ-play $\pi_{\overline{v}}$ ending in positional mode with $\mathsf{pos}(\pi_{\overline{v}}) = \overline{v}$. To construct such $\pi_{\overline{v}}$ it suffices for ∀dam to (i) follow $\overline{v}$ in positional mode (since f guides σ, the choices σ dictates to ∃ve will stay in $\overline{v}$ as well) and (ii) whenever he needs to decrement a counter for r, pick the number of remaining visits to r in $\overline{v}$ (as in the proof of (4.20)). It follows that for every f-path $\overline{v}$ from v_z to some v controlled by ∃ve, $f(\overline{v}) = \sigma(\pi_{\overline{v}})$ is defined and dictates a legal move. This proves that f is a proto-strategy, because by (4.22) there are no infinite paths in $\widetilde{\mathcal{G}}|V^\circ, v_z$ and so trivially ∃ve wins in all infinite f-paths.

We prove that for every proto-strategy f the following are equivalent:

1. ∃ve has a winning strategy σ for $\mathcal{G}, \gamma$ guided by f in $\mathcal{B}^\circ$.
2. ∃ve wins in the following game:
   (i) ∀dam picks an f-path $\overline{v}v$ starting at v_z and ending in $v \in V^\circ$;

(ii) $\exists$ve and $\forall$dam play a usual countdown game from γ but on arena restricted to $\overline{v}$ (i.e. we only update the counters and deterministically follow $\overline{v}$);

(iii) $\exists$ve wins iff the resulting configuration $(v, \mathsf{ctr}, \mathsf{cdn})$ is winning for $\exists$ve in $\mathcal{G}$.

3. $\exists$ve wins in the following game:

   (i) $\forall$dam picks an f-path $\overline{v}v$ starting at v_z and ending in $v \in V^\circ$;

   (ii) the owner of each $r \in \mathcal{D}_{\overline{v}}$ (starting from more important ranks) picks a final counter value $\mathsf{ctr}(r) < \mathsf{ctr}_z(r)$, for all other $r \in \mathcal{D} - \mathcal{D}_{\overline{v}}$ we put $\mathsf{ctr}(r) = \mathsf{ctr}_I(r)$ if r was reset in $\overline{v}$ and $\mathsf{ctr}(r) = \mathsf{ctr}_z(r)$ otherwise;

   (iii) $\exists$ve wins iff the configuration $(v, \mathsf{ctr}, \mathsf{cdn})$ is winning for $\exists$ve in $\mathcal{G}$.

The set $\mathcal{D}_{\overline{v}} \subseteq \mathcal{D}$ of nonstandard ranks in (3) contains rank r iff r appears in $\overline{v}$ and no higher rank appears after the last occurrence of r. These are the ranks for which there is a point in $\overline{v}$ when the corresponding counter is decremented and not reset to the initial value any more in $\overline{v}$. Since we are dealing with a game corresponding to a scalar formula φ, $\mathcal{D}_{\overline{v}}$ only depends on the last formula in $\overline{v}$: nonstandard r belongs to $\mathcal{D}_{\overline{v}}$ iff it is a rank of some θ which is (i) a subformula of θ_z and (ii) a superformula of the last formula in $\overline{v}$.

The implication (1) $\implies$ (2) is straightforward. Once $\forall$dam picked $\overline{v}v$, $\exists$ve simply uses σ until $\overline{v}$ is traversed. This way, $\exists$ve preserves the invariant that the play belongs to $\mathcal{B}^\circ$ (because $\overline{v} \in (V - V^\circ)^*$) and is consistent with σ (because σ is guided by f in $\mathcal{B}^\circ$ and hence $\exists$ve's positional moves dictated by f are consistent with σ). Since σ is winning, it only leads to winning configurations and therefore $\exists$ve wins (2).

To prove (2) $\implies$ (1), assume that for every f-path $\overline{v}v \in (V - V^\circ)^* V^\circ$ starting at v_z $\exists$ve has a strategy $\sigma_{\overline{v}v}$ winning in the second stage of (2) after $\forall$dam picked $\overline{v}v$ in the first stage. Our goal is to provide her with σ for (1). When during $\mathcal{B}^\circ$ she has to pick an edge, she uses:

$$\sigma(\rho) = f(\mathsf{pos}(\rho))$$

so that σ is guided by f. For choosing ordinals, observe that thanks to (4.22) for every countdown play ρ guided by f there are only finitely many f-paths extending $\mathsf{pos}(\rho)$. Hence, for every play ρ ending in $\exists$ve's choice of a counter for rank r, she takes the ordinal:

$$\sigma(\rho) = \mathsf{max}\{\sigma_{\overline{v}v}(\rho)(r) \mid \overline{v}v \in (V - V^\circ)^* V^\circ \text{ is an } f\text{-path extending } \mathsf{pos}(\rho)\}$$

which is legal, since the longer ρ is, the fewer paths extend $\mathsf{pos}(\rho)$. This way, she either wins before $\mathcal{B}^\circ$ ends, or leave it in a winning configuration, and in the later case she may continue with any winning strategy.

It remains to prove that (2) $\iff$ (3). Note that in (2), once the path $\overline{v}v$ is chosen, the only nontrivial choice of a value for r is upon the first visit to r such that no greater rank will be visited further in $\overline{v}$. If a greater rank will be visited somewhere further in $\overline{v}$, the counter for r will be reset and thus it suffices for its owner to choose the number $k < \omega$ of visits to r before the closest reset. Since $\overline{v} \in (V - V^{\circ})^*$, r was not present in the pre-modal phase and the corresponding counter must have an initial (and hence limit) value in ctr_z so this is legal. Among positions in $\overline{v}$ where r will not be reset any more, the only nontrivial choice is in the first one: in order to end the game with $\mathsf{ctr}(r) = \alpha$ it suffices to choose $\alpha + k$ where k is the number of remaining visits to r before the end of $\overline{v}$. This is legal for the same reasons as in the previous case.

It follows that these nontrivial choices in (2) are precisely the choices for $\mathcal{D}_{\overline{v}}$ and the order of the choices is precisely the (decreasing) order on $\mathcal{D}_{\overline{v}}$. This establishes an equivalence between (2) and (3), therefore completing the proof of equivalence of games (1), (2) and (3).

Let σ be a winning strategy for $\mathcal{G}, \gamma$. To complete the proof of Proposition 4.6.5 it suffices to upgrade such σ so that no formula component of shape $\langle \mathsf{a} \rangle \theta$ repeats in $\mathcal{B}^{\circ}$. Thanks to (4.21) we may apply Proposition 4.3.4 and assume that σ is guided by $\sigma^{\mathcal{B}^{\circ}}$ in $\mathcal{B}^{\circ}$. As discussed, such $\sigma^{\mathcal{B}^{\circ}}$ is a legal proto-strategy.

Enumerate all the subformulae $\langle \mathsf{a} \rangle \psi_1, ..., \langle \mathsf{a} \rangle \psi_n$ of φ of shape $\langle \mathsf{a} \rangle \theta$. We construct, by induction on $i \leq n$, a sequence $f_0, ..., f_n : V^* \to V$ of proto-strategies such that:

1. $f_0 = \sigma^{\mathcal{B}^{\circ}}$,
2. whenever $i < j$ and $\overline{v} \in V^*$ is an f_j-path, there exists an f_i-path $\overline{w} \in V^*$ ending with the same position,
3. f_i avoids repetitions of $\{\langle \mathsf{a} \rangle \psi_1, ..., \langle \mathsf{a} \rangle \psi_i\}$.

Assume we already have f_i and want to construct f_{i+1}. For every f_i-path $\overline{v} \in V^*$ ending in a visit in $\langle \mathsf{a} \rangle \psi_{i+1}$ consider the set:

$$\mathcal{H}_{\overline{v}} = \{\overline{w} \in V^* \mid \overline{w} \text{ is a } f_i\text{-path, has } \overline{v} \text{ as a prefix and ends with } \langle \mathsf{a} \rangle \psi_{i+1}\}$$

and fix some $\widehat{\overline{v}}$ maximal in $\mathcal{H}_{\overline{v}}$ ($\mathcal{H}_{\overline{v}}$ is nonempty as it contains $\overline{v}$ and must contain a maximal path because the length of paths is bounded).

Our new strategy f_{i+1} acts like f_i until the first visit in $\langle \mathsf{a} \rangle \psi_{i+1}$ and then, instead of making multiple $\xrightarrow{\mathsf{a}}$-moves for $\langle \mathsf{a} \rangle \psi_{i+1}$, immediately jumps to a *last* choice, i.e. a choice from a maximal $\widehat{\overline{v}}$ extending the current path $\overline{v}$:

$$f_{i+1}(\overline{w}) = \begin{cases} f_i(\overline{w}) & \text{if } \overline{w} \text{ does not visit } \langle \mathsf{a} \rangle \psi_{i+1}, \\ f_i(\widehat{\overline{v}} \cdot \overline{u}) & \text{if } \overline{w} = \overline{v} \cdot \overline{u} \text{ and } \overline{v} \text{ ends with the first visit in } \langle \mathsf{a} \rangle \psi_{i+1}. \end{cases}$$

Such f_{i+1} is a legal proto-strategy. Indeed, the new moves are allowed thanks to transitivity of $\xrightarrow{\mathsf{a}}$ and f_{i+1} is winning in $\widetilde{\mathcal{G}}, v_z$, because positions accessible via f_{i+1} are a subset of the ones accessible by f_i (and by (4.22) there are no infinite paths to worry about). This also implies the second property, whereas the third one follows from the fact that each $\widehat{\overline{v}}$ is maximal in $\mathcal{H}_{\overline{v}}$ (which means that no f_i-path properly extending $\widehat{\overline{v}}$ visits $\langle \mathsf{a} \rangle \psi_{i+1}$).

Since by scalarity of φ the set $\mathcal{D}_{\overline{v}}$ in the third variant of the game (3) depends only on the last formula in $\overline{v}$ and $\exists$ve wins (3) with $f = f_0$, thanks to the second property she also wins (3) with $f = f_n$. By equivalence of (3) and (1), this means that some strategy σ_{π_I} winning $\mathcal{G}, \gamma$ is guided by f_n in $\mathcal{B}^\circ$. Moreover, the third property implies that σ_{π_I} avoids repetitions of all $\langle \mathsf{a} \rangle \psi_1, ..., \langle \mathsf{a} \rangle \psi_n$ in $\mathcal{B}^\circ$, thus proving Proposition 4.6.5 and completing the proof of Theorem 4.6.2. □

4.7 Automata with Stacked Counters

A countdown automaton stores a tuple of ordinals $\alpha_1, ..., \alpha_d$ (corresponding to nonstandard ranks $r_1 < ... < r_d$) modified along the play. The modifications are restricted in two ways. First, the memory structure is hierarchical, meaning that the modifications must respect the order in the tuple. For example, decrementing α_3 requires α_1 and α_2 to be reset. Second, the counters are inherently entangled with ranks: α_3 is decremented iff the play visits rank r_3.

In this section we show that the second restriction can be dropped. That is, we introduce an equivalent model of automata where the counters are hierarchical but independent from the ranks. Such automata modify its counters $C_1, ..., C_d$ using explicit operations incorporated into the transition function (e.g. "decrement C_i", "reset C_j").

To make the counters hierarchical one could impose a syntactic requirement on the automaton: whenever C_i is modified, all C_j with $j < i$ are reset. Instead of doing that, we choose a slightly different structure which is hierarchical by design: a *stack*. The stack stores ordinals $\beta_1 ... \beta_l$ with β_l on top. Only the topmost β_l can be decremented. Apart from that, the automaton can pop the topmost counter (which results in the stack $\beta_1 ... \beta_{l-1}$) or push an ordinal β_{l+1} (obtaining $\beta_1 ... \beta_l \beta_{l+1}$). The intuition is that a configuration of a countdown automaton with ordinals $\alpha_1 ... \alpha_d$ is represented as a stack $\alpha_d ... \alpha_1$ (i.e. with values for greater ranks being closer to the bottom of the stack). To avoid different players manipulating the same counter, we will assume that each β_i on the stack comes with an owner $P_i \in \{\exists, \forall\}$ controlling it. Thus, formally the stack content will be of shape $(P_1, \beta_1) ... (P_l, \beta_l)$.

Definition 4.7.1. An *automaton with a stack of counters* $\mathcal{A}$ consists of:

- a finite set of states $Q = Q_\exists \sqcup Q_\forall$ divided between two players;

- an initial state $q_I \in Q$;
- a transition function:

$$\delta : Q \to \mathcal{P}(Q \sqcup \mathsf{Lit} \sqcup \mathsf{Var}) \sqcup (\mathsf{Act} \times Q) \sqcup (\mathsf{Ops} \times Q)$$

  where the set:

$$\mathsf{Ops} = \{\mathsf{push}(P, \alpha) \mid P \in \{\exists, \forall\}, \alpha \in \mathsf{Ord}\} \sqcup \{\mathsf{pop}, \mathsf{dec}\}$$

  is called *stack operations*;
- a ranks function $\mathsf{rank} : Q \to \mathcal{R}$.

That is, $\mathcal{A}$ is a parity automaton (as defined in Subsection 2.4.3) except that it can additionally perform *stack-transitions*, i.e. transitions labelled with stack operations.

The semantics of an automaton with a stack of counters $\mathcal{A}$ is defined by a parity game similarly to parity automata. The only difference is that now configurations of the game contain a *stack of ordinal counters* updated along the play.

Definition 4.7.2. Fix a model $\mathcal{M}$. A configuration of an automaton consists of a point $\mathsf{m} \in M$, a state $q \in Q$ and a stack content $\Upsilon = (P_1, \alpha_1)...(P_l, \alpha_l)$ where $l < \omega$, $\alpha_i \in \mathsf{Ord}$ and $P_i \in \{\exists, \forall\}$ is the owner of the i-th counter. The operations $\mathsf{op} \in \mathsf{Ops}$ have the following meaning:

- $\mathsf{push}(P, \alpha)$ puts (P, α) at the top of the stack;
- pop removes the topmost counter from the stack;
- dec, which is the only nondeterministic operation, means that we look at the topmost counter (P, α), player P picks some $\beta < \alpha$ and the new stack is the old one with the topmost counter (P, α) replaced with (P, β)

If at any moment one of the players has to decrement 0, he or she looses immediately. It will be less important what happens when we try to pop or decrement an empty stack: assume by convention that in such a situation nothing happens. Formally:

$$\mathsf{Conf} = M \times Q \sqcup \times(\{\exists, \forall\} \times \mathsf{Ord})^*.$$

There are no moves from configurations of shape $\gamma = (\mathsf{m}, z, \Upsilon)$ with z being a literal or a variable. From $(\mathsf{m}, q, \Upsilon)$ with $\Upsilon = (P_1, \alpha_1)...(P_l, \alpha_l)$:

- if $\delta(q) \subseteq Q \sqcup \mathsf{Lit} \sqcup \mathsf{Var}$, outgoing moves lead to:

$$\{(\mathsf{m}, z, \Upsilon) \mid z \in \delta(q)\},$$

- if $\delta(q) = (\mathsf{a}, p)$, outgoing moves lead to:
$$\{(\mathsf{n}, p, \Upsilon) \mid \mathsf{m} \xrightarrow{\mathsf{a}} \mathsf{n}\}.$$
- if $\delta(q) = (\mathsf{op}, p)$ define the set $\mathcal{H}$ of possible new contents of the stack:
  - if $\mathsf{op} = \mathsf{push}(P, \alpha)$:
$$\mathcal{H} = \{(P_1, \alpha_1)...(P_l, \alpha_l)(P, \alpha)\}$$
  - if $\mathsf{op} = \mathsf{pop}$:
$$\mathcal{H} = \{(P_1, \alpha_1)...(P_{l-1}, \alpha_{l-1})\}$$
    (or $\mathcal{H} = \{\varepsilon\}$ if $\Upsilon = \varepsilon$)
  - if $\mathsf{op} = \mathsf{dec}$, then:
$$\mathcal{H} = \{(P_1, \alpha_1)...(P_{l-1}, \alpha_{l-1})(P_l, \beta) \mid \beta < \alpha_l\}$$

  and the outgoing moves lead to:
$$\{(\mathsf{m}, p, \Upsilon') \mid \Upsilon' \in \mathcal{H}\}.$$

The ownership of positions is the same as with simple parity automata except for configurations of shape $(\mathsf{m}, q, \Upsilon)$ with $\delta(q) = (p, \mathsf{op})$ which belong to player P_l from the top of the stack Υ (or any player if the stack is empty). The rank function $\mathsf{rank} : \mathsf{Conf} \to \mathcal{R}$ is inherited from the states of $\mathcal{A}$. The default initial configuration contains the initial state q_I and an empty stack.

The following example demonstrates that arbitrary automata with stacked counters have expressive power strictly greater than countdown automata from Section 4.4. This is because the stack structure can be exploited in order to define context-free languages.

Example 4.7.3. Consider the following automaton:

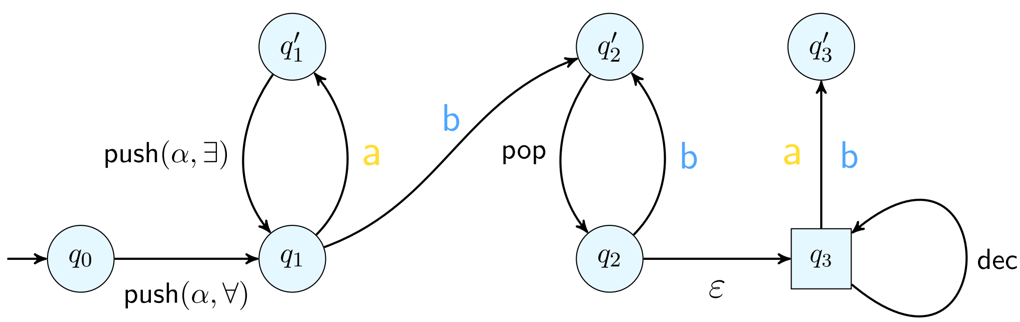

The only state q_3 controlled by $\forall$dam is depicted with a square and all the other states are controlled by $\exists$ve and depicted with circles; q_0 is initial. There is only one rank, belonging to $\exists$ve, and the value $\alpha \in \mathsf{Ord}$ is irrelevant. For clarity of presentation we allow for multiple transitions other than ε-ones to originate in the same state (e.g. in q_1 $\exists$ve can choose either an $\xrightarrow{\mathsf{a}}$- or $\xrightarrow{\mathsf{b}}$- transition). Such a more liberal syntax is equivalent to our more rigorous one: to fit to the original syntax it suffices to introduce additional states (e.g. q_1^{a} and q_1^{b} such that in q_1 first $\exists$ve chooses an ε-move to one of them and then from q_1^{c} fires an $\xrightarrow{\mathsf{c}}$-transition).

We view $\mathcal{A}$ as a device accepting finite words over the alphabet $\Gamma = \{\mathsf{a}, \mathsf{b}\}$. A word $\mathsf{a}_1...\mathsf{a}_l \in \Gamma^*$ of length l is encoded as a model with edges labelled with Γ. The encoding is the same as in Example 4.2.5 except that the words are finite so in place of ω we take $\{0, ..., l\}$ to be the universe.

We claim that for every word $w \in \Gamma^*$:

$$\mathcal{A} \text{ accepts } w \iff w = \mathsf{a}^n\mathsf{b}^n \text{ for some } n < \omega.$$

The corresponding semantic game can be divided into 5 stages:

1. The stack is initialized with $(\forall, \alpha)$ and the automaton moves to q_1.

2. As long as there are no b's, the automaton stays in q_1 and for each a one counter $(\exists, \alpha)$ is pushed to the stack.

3. Upon reaching the first b the state changes to q_2 and each b pops one counter from the stack (if at this stage a is visited, $\exists$ve gets stuck).

4. At some point $\exists$ve decides to move from q_2 to q_3. Note that she looses immediately in q_3' and if the current position is not the last position l of the word $w = \mathsf{a}_1...\mathsf{a}_l$ then $\forall$dam can move from q_3 to q_3' and win. Therefore, without loosing generality we may assume that she decides for this ε-move only at the end of the word.

5. The state q_3 is reached at the end of the word and $\forall$dam has no choice but to loop indefinitely decrementing the topmost counter on the stack (meaning that he wins if the topmost counter belongs to $\exists$ve or the stack is empty).

The only way $\exists$ve can win the game is by forcing $\forall$dam to decrement his counter indefinitely until he gets stuck. For that, she needs to reach (5) with $\forall$dam's counter on the top of the stack. Since the only time when his counter is pushed to the stack is at the beginning of the game, it follows that $\exists$ve wins iff w is of shape $\mathsf{a}^n\mathsf{b}^n$ for some $n < \omega$.

On the other hand, the language of $\mathcal{A}$ cannot be recognized by a countdown automaton. Since countdown automata are equivalent to $\mu^{<\infty}$-ML and in finite models all fixpoints are reached after finitely many steps, it follows that countdown automata only define regular languages of *finite* words, i.e.

languages definable in μ-ML. The language $\mathcal{L}(\mathcal{A}) = \{\mathsf{a}^n\mathsf{b}^n \mid n < \omega\}$ is a prototypical example of a language that is not regular (as can be shown using standard techniques) and hence cannot be defined in $\mu^{<\infty}$-ML.

The above example exploits the fact that the stack can have unbounded height. However, if we require that every automaton has a finite bound k on the maximal height of its stack, then the resulting model turns out to be equivalent to countdown automata. From now on we tacitly assume that every automaton with stacked counters comes with a bound k on the height of its stack and never pops or decrements an empty stack. In Subsection 4.7 we will see that such k is a good measure of expressive power. Before that, let us prove that countdown automata and automata with stacked counters are indeed equivalent.

Theorem 4.7.4. *Countdown automata and automata with stacked counters define the same languages.*

Proof. The direction from countdown automata to automata with stacked counters is rather straightforward and so we only give a sketch. Observe that using a stack we may represent the configuration of any given countdown automaton $\mathcal{A}$. The counter values:

$$\mathsf{ctr} : \mathcal{D} \to \mathsf{Ord}$$

for $\mathcal{D} = \{r_1 < ... < r_d\}$ are represented as a stack:

$$\Upsilon = (P_d, \mathsf{ctr}(r_d))...(P_1, \mathsf{ctr}(r_1))$$

where P_i is the owner of r_i. In $\mathcal{A}$, whenever the counter for rank r_i is decremented, counters for all the less important ranks are reset. Thus, we may pop $(P_{i-1}, \mathsf{ctr}(r_{i-1}))...(P_1, \mathsf{ctr}(r_1))$ from the stack, ask the owner P_i of r_i to decrement the counter representing $\mathsf{ctr}(r_i)$, and then push initial values $(P_{i-1}, \mathsf{ctr}_I(r_{i-1})...(P_1, \mathsf{ctr}_I(r_1))$ to the stack. Such simulation requires a stack of height equal to the number of nonstandard ranks $|\mathcal{D}|$ of the original automaton.

For the more demanding direction, take an automaton $\mathcal{A} = (Q, q_I, \delta, \mathsf{rank})$ with stacked counters and stack bounded by k. For simplicity, assume that there is only one ordinal κ_I that can be pushed to the stack (the general case is not harder and we will comment on that at the end of the proof). Given a stack content $\Upsilon = (P_1, \alpha_1)...(P_l, \alpha)$ of $\mathcal{A}$, we call the sequence of ownerships $P_1...P_l$ the *shape* of Υ. Let us start with the following lemma.

Lemma 4.7.5. *The automaton $\mathcal{A}$ can be always modified so that:*

1. *$\mathcal{A}$ remembers the current shape of the stack in its state.*

2. *Denote the states in which the stack has height i by Q^i and their ranks $\mathsf{rank}[Q^i]$ by $\mathcal{R}^i$. For all $i < j$:*

$$\mathcal{R}^i > \mathcal{R}^j$$

   where by $X > Y$ we mean that $x > y$ for all $x \in X$ and $y \in Y$.

   In words: the higher the stack, the lower the rank.

Proof. The first item (1) is immediate. The shape of the stack is a finite information that can be easily updated. We assume an additional property (which will be lost in further modifications of the automaton): if q is a target of a dec-transition $p \overset{\mathsf{dec}}{\to} q$ then no other transition has q as its target and q has the irrelevant lowest rank 0. This can be easily obtained by adding more states: replace every $p \overset{\mathsf{dec}}{\to} q$ with $p \overset{\mathsf{dec}}{\to} p' \overset{\varepsilon}{\to} q$ with p being a fresh state of rank 0.

Assuming the above we massage the automaton so that it satisfies the second item (2). We do so by defining a new rank function $\mathsf{rank}' : Q \to \mathcal{R}'$. Take $k + 1$ copies:

$$\mathcal{R}' = \mathcal{R}^0 \sqcup ... \sqcup \mathcal{R}^k$$

of the original set $\mathcal{R}$ of ranks, one for each stack height. The ownership in $\mathcal{R}'$ is inherited from $\mathcal{R}$. The new order $<'$ is the same as $<$ inside each copy and orders the copies descendingly:

$$\mathcal{R}^0 > ... > \mathcal{R}^k.$$

A naïve solution would be to assign the i-th copy r^i of r to q iff $\mathsf{rank}(q) = r$ and $q \in Q^i$. In plays π where the height of the stack is eventually constant, this modified rank function would then give the same winner as the original rank. This is because all but finitely many states in such π belong to the same Q^i and hence from some point in π the new rank would be equal r^i iff the original equals r.

However, an issue could arise if the height of the stack changes infinitely often in π. Whenever $i < j$, the copies r_0^i and r_1^j of ranks $r_0 < r_1$ have swapped order $r_1^i >' r_0^j$. This could change the winner of the play, e.g. if r_0 is only visited with a stack of height i and r_1 only with height j and these are the only ranks seen infinitely often in π. To address that, we assume that for each $0 \le i \le l$ the automaton stores in its state the rank $\mathsf{since}(i) \in \mathcal{R}$:

$$\mathsf{since}(q, i) = \begin{array}{l}\text{the greatest value of } \mathsf{rank} \text{ visited after the last} \\ \text{time when the stack had height } i.\end{array}$$

We want the new automaton $\mathcal{A}'$ to be the same as $\mathcal{A}$ except that upon each $p \overset{\mathsf{pop}}{\to} q$ of the stack from height $i+1$ to i, the new automaton additionally passes through the i-th copy $(\mathsf{since}(p,i))^i \in \mathcal{R}'$ of $\mathsf{since}(p,i)$. Formally, the new rank function $\mathsf{rank}' : Q \to \mathcal{R}'$ is given as:

$$\mathsf{rank}'(q) = \begin{cases} (\mathsf{since}(p,i))^i & \text{if } q \text{ is a target of some } p \overset{\mathsf{pop}}{\to} q, \\ r^i & \text{otherwise with } \mathsf{rank}(q) = r \end{cases}$$

for every $q \in Q_i$. The resulting automaton $\mathcal{A}'$ has both the desired properties (1) and (2) so it remains to prove that it is equivalent to the original $\mathcal{A}$. Fix a model $\mathcal{M}$ and a valuation val. The arenas of the semantic games $\mathcal{G}^{\mathsf{val}}(\mathcal{A})$ and $\mathcal{G}^{\mathsf{val}}(\mathcal{A}')$ are identical, so we only need to check that the winner of any infinite play π is the same with rank and rank'. Let r and r' be the most important ranks appearing infinitely often in $\mathsf{rank}[\pi]$ and $\mathsf{rank}'[\pi]$, respectively. We claim that r' is the i-th copy of r:

$$r' = r^i \tag{4.23}$$

which implies equivalence of the two games because r^i has the same owner as r.

Let i be the smallest stack height present infinitely often in π. The case when the height of the stack is eventually constant is already discussed: if this happens then after some finite prefix of π rank' always returns the i-th copy of the rank returned by rank. Otherwise, π can be decomposed into:

$$\pi = \pi_0 \pi_1 \pi_2 ...$$

such that (i) π_0 contains all the configurations with rank greater than r and stack smaller than i and (ii) for every $0 < n < \omega$, the fragment π_n visits r and ends with a pop from height $i+1$ to i. It follows that the most important rank in π_n, denoted $r_n = \mathsf{max}(\mathsf{rank}[\pi_n])$, equals r for all n. It suffices to prove that $r'_n = \mathsf{max}(\mathsf{rank}'[\pi_n])$ equals r^i. Observe that at the beginning of π_n the stack has height i, so $\mathsf{since}(p,i) \leq r$ for all p in π_n. Any q in π_n is a target of some $p \overset{\mathsf{pop}}{\to} q$ iff it is preceded by p also belonging to π_n. It follows that no state q visited in π_n could have $\mathsf{rank}'(q) > r^i$, as this would require either a stack smaller than i or $\mathsf{rank}(q) > r$, both of which contradict (i). Hence, it suffices to show that r^i does appear in $\mathsf{rank}'[\pi_n]$.

By (ii), some state s in π_n has $\mathsf{rank}(s) = r$. By definition of 0, $0 < r$ and hence s is not a target of a pop-transition. Therefore, if s is visited with stack of height i then $\mathsf{rank}'(s) = r^i$ and we are done. Otherwise, s is visited with stack of height $j > i$. Let π'_n be the prefix of π_n ending with the first pop of the stack from height $i+1$ to i after a visit to s. The fragment π'_n ends with a transition $p \overset{\mathsf{pop}}{\to} q$ for some p and q. It follows that the value $\mathsf{since}(p,i)$ equals r and consequently $\mathsf{rank}'(q) = r^i$. $\square$

With the above lemma, we prove Theorem 4.7.4. We assume that the automaton with stacked counters $\mathcal{A} = (Q, q_I, \delta, \mathsf{rank})$ has properties:

1. $\mathcal{A}$ remembers the shape of the stack in its state.
2. Whenever $p \in Q^i - Q^{\mathsf{Ops}}$ and $q \in Q^j - Q^{\mathsf{Ops}}$ with $i < j$:
$$\mathsf{rank}(p) > \mathsf{rank}(q)$$
and:
$$\mathsf{rank}(q) = 0$$
whenever $q \in Q^{\mathsf{Ops}}$, where $Q^{\mathsf{Ops}} = \{q \mid \delta(q) \in \mathsf{Ops} \times Q\}$ is the set of all the states that are sources of stack-transitions.

   That is: sources of stack-transitions have rank 0; and for all the other states the higher the stack the lower the rank.
3. $\mathcal{A}$ never fires two stack-operations $p \overset{\mathsf{op}}{\to} q \overset{\mathsf{op}'}{\to} s$ in a row.

The first property (1) is (1) from Lemma 4.7.5. The second (2) and third (3) can be easily obtained from (2) in the lemma by adding more states to $\mathcal{A}$, without breaking (1): replace every $p \overset{\mathsf{op}}{\to} q$ with $p \overset{\mathsf{op}}{\to} p' \overset{\varepsilon}{\to} q$ for a fresh p', assign $\mathsf{rank}(p)$ to p' and 0 to p.

We define a countdown automaton $\mathcal{A}' = (Q, \delta', q_I, \mathsf{rank}', \mathsf{ctr}_I)$ equivalent to $\mathcal{A}$. The new δ' is obtained from δ by replacing all the stack-transitions with ordinary ε-transitions. The codomain $\mathcal{R}'$ of $\mathsf{rank}' : Q \to \mathcal{R}'$ extends $\mathcal{R}$:

$$\mathcal{R}' = \mathcal{R} \sqcup \mathcal{D}$$

by nonstandard:

$$\mathcal{D} = \{t^i_\exists, t^i_\forall \mid 1 \leq i \leq k\}$$

with each t^i_P put between $\mathcal{R}^{i-1}$ and $\mathcal{R}^i$:

$$\mathcal{R}^0 >' \{t^1_\exists, t^1_\forall\} >' \mathcal{R}^1 >' ... >' \{t^k_\exists, t^k_\forall\} >' \mathcal{R}^k$$

(the order between $t^i_\exists$ and $t^i_\forall$ does not matter). The new rank function is:

$$\mathsf{rank}'(p) = \begin{cases} t^l_{P_l} & \text{if } p \overset{\mathsf{dec}}{\to} q \text{ for some } q\text{, with stack of shape } P_1...P_l; \\ \mathsf{rank}(p) & \text{if } p \overset{\mathsf{op}}{\to} q \text{ for some } q \text{ and } \mathsf{op} \neq \mathsf{dec}; \\ \mathsf{rank}(q) & \text{otherwise} \end{cases}$$

for every p. That is, for every stack-transition $p \overset{\mathsf{op}}{\to} q$ we replace the irrelevant $\mathsf{rank}(p) = 0$ with (i) appropriate t^i_P if $\mathsf{op} = \mathsf{dec}$ or (ii) rank of the target $\mathsf{rank}(q)$ if $\mathsf{op} \neq \mathsf{dec}$. On all the states which are not sources of stack-transitions, rank' and rank are the same. The initial $\mathsf{ctr}_I : \mathcal{D} \to \mathsf{Ord}$ is the constant function equal κ_I.

The semantic games $\mathcal{G} = \mathcal{G}^{\mathsf{val}}(\mathcal{A})$ and $\mathcal{G}' = \mathcal{G}^{\mathsf{val}}(\mathcal{A}')$ for a fixed model $\mathcal{M}$ and valuation val are almost the same. For a stack content $\Upsilon = (P_1, \alpha_1)...(P_l, \alpha_l)$, define a counter assignment $\mathsf{ctr}_\Upsilon : \mathcal{D} \to \mathsf{Ord}$:

$$\mathsf{ctr}_\Upsilon(r) = \begin{cases} \alpha_i & \text{if } r = t^i_{P_i} \text{ with } i \leq l \\ \kappa_I & \text{otherwise.} \end{cases}$$

It follows that:

$$\mathsf{ctr}_\Upsilon(r) \neq \kappa_I \quad \Longrightarrow \quad r \in \{t^1_{P_1}, ..., t^l_{P_l}\} \tag{4.24}$$

for every $r \in \mathcal{D}$. Consider the relation $\sim$ between all the configurations of $\mathcal{G}$ and countdown configurations of $\mathcal{G}'$:

$$(\mathsf{m}, p, \Upsilon) \sim (\mathsf{m}, p, \mathsf{ctr}_\Upsilon, \mathsf{cdn})$$

for every point m, stack content Υ and state p and consistent with the shape of Υ. We claim that $\sim$ is a bisimulation between $\mathcal{G}$ and $\mathcal{G}'$, up to identifying moves $\gamma_0 \to \gamma_1$ in $\mathcal{G}$ with pairs of moves $\gamma'_0 \to \gamma'_{\frac{1}{2}} \to \gamma'_1$ in $\mathcal{G}'$. That is, whenever $\gamma_0 \sim \gamma'_0$:

- (forth) for every move $\gamma_0 \to \gamma_1$ in $\mathcal{G}$ then there are $\gamma'_0 \to \gamma'_{\frac{1}{2}} \to \gamma'_1$ in $\mathcal{G}'$ with $\gamma_1 \sim \gamma'_1$,
- (back) for every $\gamma'_0 \to \gamma'_{\frac{1}{2}} \to \gamma'_1$ in $\mathcal{G}'$ then there is $\gamma_0 \to \gamma_1$ in $\mathcal{G}$ with $\gamma_1 \sim \gamma'_1$.

In both cases, only one move in $\mathcal{G}'$ is nondeterministic and the same player makes the nontrivial choice in $\mathcal{G}$ and $\mathcal{G}'$. Moreover, $\sim$ preserves victory in a sense that will be explained later.

Towards the back and forth conditions, take configurations:

$$\gamma_0 = (\mathsf{m}, p, \Upsilon) \sim (\mathsf{m}, p, \mathsf{ctr}_\Upsilon, \mathsf{cdn}) = \gamma'_0$$

as above. Since the stack has height l, the next visited state after p must have stack height at least $l-1$. It follows that if $p \in Q^{\mathsf{Ops}}$ then $\mathsf{rank}'(p) \in \mathcal{R}^{l-1} \sqcup \{t^l_{P_l}\} \sqcup \mathcal{R}^{l+1}$ (or $\mathsf{rank}'(p) \in \mathcal{R}^{l+1}$ if $l = 0$) and otherwise $\mathsf{rank}'(p) \in \mathcal{R}^l$. In the presence of (4.24) this means that in the upcoming counter update $\gamma'_0 \to \gamma'_{\frac{1}{2}}$ in $\mathcal{G}'$ the only counter value that can change is the one for $t^l_{P_l}$, and it only changes if $\mathsf{rank}'(p) \geq' t^l_{P_l}$.

- If $\delta(p) \notin Q^{\mathsf{Ops}}$ is not a source of a stack-transition then $\mathsf{rank}'(p) = \mathsf{rank}(p) \in \mathcal{R}^l$. Hence, the update $\gamma'_0 \to \gamma'_{\frac{1}{2}}$ does not change the counters at all: deterministically $\gamma'_{\frac{1}{2}} = (\mathsf{m}, p, \mathsf{ctr}_\Upsilon, \mathsf{psn})$. By definition of δ', the owner of p has a move from $\gamma'_{\frac{1}{2}}$ to $\gamma'_1 = (\mathsf{n}, q, \mathsf{ctr}_\Upsilon, \mathsf{cdn})$ in $\mathcal{G}'$ iff (s)he has a move from γ_0 to $\gamma_1 = (\mathsf{n}, q, \Upsilon)$ in $\mathcal{G}$. Clearly, $\gamma_1 \sim \gamma'_1$.

- On the other hand, if $\delta(p) = (\mathsf{op}, q)$ then the choice of $\mathsf{ctr}_{\frac{1}{2}}$ in $\gamma'_0 \to \gamma'_{\frac{1}{2}}$ may be nontrivial. Consider cases:
  - If $\mathsf{op} = \mathsf{push}(P, \kappa_I)$ then $\mathsf{rank}'(p) = \mathsf{rank}(q) \in \mathcal{R}^{l+1}$ is standard and by (4.24) smaller than all r for which $\mathsf{ctr}_\Upsilon(r) \neq \kappa_I$. Thus, deterministically:
  $$\mathsf{ctr}_{\frac{1}{2}} = \mathsf{ctr}_\Upsilon = \mathsf{ctr}_{\Upsilon'}$$
  where $\Upsilon' = (P_1, \alpha_1)...(P_l, \alpha_l)(P, \kappa_I)$.
  - If $\mathsf{op} = \mathsf{pop}$ then $\mathsf{rank}'(p) = \mathsf{rank}(p) \in \mathcal{R}^{l-1}$ is also standard and smaller than all $t^i_{P_i}$ with $i < l$ but greater than $t^l_{P_l}$ (whose value is therefore reset). It follows that deterministically:
  $$\mathsf{ctr}_{\frac{1}{2}} = \mathsf{ctr}_\Upsilon[t^l_{P_l} \mapsto \kappa_I] = \mathsf{ctr}_{\Upsilon'}$$
  where $\Upsilon' = (P_1, \alpha_1)...(P_{l-1}, \alpha_{l-1})$.
  - If $\mathsf{op} = \mathsf{dec}$ then $\mathsf{rank}'(p) = t^l_{P_l}$ is nonstandard. Thus, the player P_l has to decrement $\mathsf{ctr}_\Upsilon(t^l_{P_l}) = \alpha_l$ by choosing $\beta < \alpha_l$, as in $\mathcal{G}$.

  In either case, P_l can choose $\mathsf{ctr}_{\frac{1}{2}}$ in $\mathcal{G}'$ iff (s)he can choose Υ' in $\mathcal{G}'$ such that $\mathsf{ctr}_{\frac{1}{2}} = \mathsf{ctr}_{\Upsilon'}$. Such choice of $\mathsf{ctr}_{\Upsilon'}$ in $(\mathsf{m}, p, \mathsf{ctr}_{\Upsilon'}, \mathsf{psn}) = \gamma'_{\frac{1}{2}}$ is followed by a deterministic ε-move leading to $\gamma'_1 = (\mathsf{m}, q, \mathsf{ctr}_{\Upsilon'}, \mathsf{cdn})$. Again, denoting $\gamma_1 = (\mathsf{m}, q, \Upsilon')$ we have $\gamma_1 \sim \gamma'_1$.

This completes the proof of the back and forth conditions. Let us show that $\sim$ is victory-preserving in a suitable sense. Take infinite plays $\pi = \gamma_0\gamma_1\gamma_2...$ in $\mathcal{G}$ and $\pi' = \gamma'_0\gamma'_{\frac{1}{2}}\gamma'_1\gamma'_{1\frac{1}{2}}\gamma_2...$ in $\mathcal{G}'$ with $\gamma_i \sim \gamma'_i$ for all integer i. We claim that the most important ranks r and r' appearing infinitely often in $\mathsf{rank}[\pi]$ and $\mathsf{rank}'[\pi']$, respectively, are the same. Every countdown configuration in a countdown game has the same rank as the next positional configuration. Hence, to determine r' it suffices to look at the configurations $\gamma'_1, \gamma'_2...$ with integer indices in π'. We only have $\mathsf{rank}(q) \neq \mathsf{rank}'(q)$ if $q \in Q^{\mathsf{Ops}}$ so it suffices to prove that the visits to Q^{Ops} can be ignored while determining r and r'. This is clearly true for r, because $q \in Q^{\mathsf{Ops}}$ implies $\mathsf{rank}(q) = 0$. Assume $q \overset{\mathsf{op}}{\to} p$ for some p. If $\mathsf{op} \neq \mathsf{dec}$ then $\mathsf{rank}'(q) = \mathsf{rank}(p) = \mathsf{rank}'(p)$. Since every visit to q is followed by a visit to p, we may ignore visits to q while computing r'. Otherwise $\mathsf{op} = \mathsf{dec}$ and $\mathsf{rank}'(q)$ is nonstandard. As such, it cannot be seen infinitely often without infinitely many visits to a more important rank, meaning that $\mathsf{rank}'(q) <' r'$. It follows that visits to Q^{Ops} can be indeed ignored when it comes to computing r' and so $r = r'$.

The above observations about $\sim$ can be used to translate strategies between $\mathcal{G}$ and $\mathcal{G}'$. To be fully formal, we may apply the Decomposition Lemma 2.2.5 with S_+ equal to $\sim$ and $S = \emptyset$. The above discussion can be summarized as exit-equivalence $\mathcal{G}|\mathcal{S}_+, \gamma \bowtie_{S_+} \mathcal{G}'|\mathcal{S}'_+, \gamma'$ for all $\gamma \sim \gamma'$ and victory-dominance of $\sim$ between $\mathcal{G}|\mathcal{S}_+$ and $\mathcal{G}'|\mathcal{S}'_+$. □

Remark 4.7.6. The above proof assumes that the countdown always starts from a fixed κ_I, but this assumption can be easily dropped. With several possible initial values $\kappa_1, ..., \kappa_n$ it suffices for $\mathcal{A}$ to remember, for each counter (P_i, α_i) on the stack, the initial κ_j with which it was pushed. Then, in place of the nonstandard ranks $t^i_\exists$ and $t^i_\forall$ the constructed countdown automaton $\mathcal{A}'$ has $t^i_{\exists,\kappa_1}, ..., t^i_{\exists,\kappa_n}$ and $t^i_{\forall,\kappa_1}, ..., t^i_{\forall,\kappa_n}$. Upon each dec-transition of $\mathcal{A}$, the new $\mathcal{A}'$ visits appropriate $t^i_{P_i,\kappa_j}$ (instead of $t^i_{P_i}$).

4.7.1 Stack Bound = Nesting = Countdown Depth

The maximal height k of the stack reflects how many ordinals need to be remembered at once. The same automaton with stack of height one can keep either $(\exists, \alpha)$ or $(\forall, \alpha)$ in its memory. Hence, it can describe boolean combinations of boundedness and unboundedness conditions such as:

"There are arbitrarily long $\xrightarrow{\mathsf{a}}$-paths *and* no arbitrarily long $\xrightarrow{\mathsf{b}}$-paths."

It can be easily shown that any countdown formula describing such property needs to have at least two countdown operators: ν^ω and μ^ω. The operators need not be nested, though, as witnessed by the formula:

$$(\nu^\omega x.\langle \mathsf{a} \rangle x) \wedge (\mu^\omega y.[\mathsf{b}]y)$$

expressing the property. Nesting of the countdown operators (being the maximal length k of a chain of its distinct subformulae $\theta_1, ..., \theta_k$, each beginning with a countdown operator) is a more fine-grained parameter than their overall number. For instance, the above formula has nesting 1 despite using two countdown operators. It is the nesting of the operators, not their number, that corresponds to the number of ordinals remembered at once.

Similarly, a naïve approach to the complexity of countdown automata would be to restrict the number of non-standard ranks. However, this does not faithfully capture the intuition that at most k numbers need to be remembered at once. Every automaton recognizing the discussed language has at least one nonstandard rank for each player. The corresponding counters, thought, need not be used simultaneously. To capture that, we introduce the following parameter.

Definition 4.7.7. Consider a countdown $\mathcal{A}$ automaton with nonstandard ranks $\mathcal{D}$. The *countdown depth* of $\mathcal{A}$ is the least $k < \omega$ for which $\mathcal{D}$ can be divided into disjoint sets $\mathcal{D}_1, ..., \mathcal{D}_k$ such that whenever a path π in the automaton visits two different ranks $r, r' \in \mathcal{D}_i$ from the same set, then between these visits π passes through a (possibly standard) rank r'' greater than both r and r'.

The definition implies that at most one rank in each $\mathcal{D}_i$ has a counter with non-initial value. Consequently, at most k counters in total have non-initial values.

The next theorem says that the three measures of complexity coincide. Countdown formulae with nesting of the countdown operators at most k, countdown automata of countdown depth k and automata with stacked counters with stack height k are equivalent.

Theorem 4.7.8. *For every $k < \omega$ and language L of modal models, it is equivalent for L to be recognized by:*

1. *a countdown sentence with nesting of the countdown operators k;*

2. *a countdown automaton of countdown depth k;*

3. *an automaton with stacked counters with stack height bounded by k.*

Proof. In the proof we provide translations between automata of both types and logic. In each case, the translation maps a language-defining object (a sentence or an automaton) with respective parameter k to one with parameter at most k.

(1) $\implies$ (2)
Take a sentence $\varphi \in \mu^{<\infty}\text{-}\mathsf{ML}$ with nesting of the countdown operators equal k. We claim that φ seen as an automaton $\mathcal{A}_\varphi$ has countdown depth k. To show that, we need to decompose nonstandard ranks $\mathcal{D}$ of $\mathcal{A}_\varphi$ into disjoint $\mathcal{D}_1, ..., \mathcal{D}_k$ satisfying the definition. Let $Y \subseteq \mathsf{SubFor}(\varphi)$ be all the subformulae of φ that begin with a countdown or fixpoint operator and $X \subseteq Y$ be the ones that begin with a countdown (but not fixpoint) operator. Assume that the rank function $\mathsf{rank} : \mathsf{SubFor}(\varphi) \to \mathcal{R}$ assigns different nonstandard ranks whenever possible: ψ and ψ' have the same nonstandard rank only if they are both immediate subformulae of the same $\theta \in X$. This allows to identify nonstandard ranks with elements of X (each $\xi \in X$ corresponds to the rank of its immediate subformulae). Hence, it suffices to decompose X into disjoint $X_1, ..., X_k$ such that for every distinct ξ, ξ' belonging to the same X_i any path between its respective immediate subformulae ψ and ψ' passes through a rank greater than both $\mathsf{rank}(\psi)$ and $\mathsf{rank}(\psi')$.

Denote the subformula order on $\mathsf{SubFor}(\varphi)$ by $\preceq$, i.e. $\psi \preceq \psi'$ iff $\psi \in \mathsf{SubFor}(\psi')$. This $\preceq$ is a linear order on superformulae of any given formula. We claim that in $\mathcal{A}_\varphi$, for every path π leading from ψ to ψ' either:

(i) $\psi' \preceq \psi$ or

(ii) there is $\theta \in Y$ and its immediate subformula $\chi \prec \theta$ such that $\psi \prec \theta$, $\psi' \preceq \chi$ and π passes through χ.

We prove the above by induction on the length of π. For $|\pi| = 1$ the claim is obvious with $\psi = \psi'$. Assume it is true for paths of length n and consider π of length $n+1$. Decompose $\pi = \pi_0\psi''$ with π_0 of length n leading from ψ to ψ'. If $\psi'' \preceq \psi'$ then the claim follows immediately from the induction hypothesis applied to π_0. Otherwise, ψ' must be a variable $y \in \mathsf{BoundVar}(\varphi)$ bound in some its superformula $\theta' \in Y$ having $\psi'' \prec \theta'$ as an immediate subformula.

Note that if ψ and θ' are $\preceq$-comparable then we are done. On the one hand, $\theta' \preceq \psi$ implies $\psi'' \preceq \theta' \preceq \psi$ so π satisfies the first item (i) of the claim. On the other hand, $\psi \prec \theta'$ implies that both ψ and ψ'' are proper subformulae of $\theta' \in Y$ and π ends in ψ''. Thus, π satisfies the second item (ii) witnessed by θ' in place of θ and ψ'' in place of χ.

Applying the induction hypothesis to π_0 we get two cases.

(a) If $y \preceq \psi$, then ψ and θ' are both superformulae of y. Hence, they are $\preceq$-comparable and we are done.

(b) Otherwise assume $\theta \in Y$ with an immediate subformula χ such that $\psi \prec \theta$, $y \preceq \chi$ and π_0 passes through χ. Since both θ and θ' are superformulae of y, they are $\preceq$-comparable. If $\theta \prec \theta'$ then we get $\preceq$-comparability of ψ and θ': $\psi \prec \theta \prec \theta'$. In the remaining case $\theta' \preceq \theta$ implies $\psi'' \prec \theta' \preceq \theta$. This means that π satisfies the second item (ii) of the claim, witnessed by the same θ and χ.

This completes the proof of the claim about paths in $\mathcal{A}_\varphi$, so let us use it. If $\xi, \xi' \in X$ are $\preceq$-incomparable then any two their respective immediate subformulae ψ and ψ' are $\preceq$-incomparable as well. Hence, the claim implies that every path π from ψ to ψ' must pass through an immediate subformula χ of some $\theta \in Y$ with $\psi \prec \theta$ and $\psi' \preceq \chi$. Since ξ is the immediate superformula of ψ, every proper superformula of ψ is (not necessarily proper) superformula of ξ. In particular, $\psi \prec \theta$ implies $\xi \preceq \theta$. Symmetrically, we get $\xi' \preceq \theta$. By $\preceq$-incomparability of ξ and ξ' the inequalities are strict: $\xi, \xi' \prec \theta$. This means that $\mathsf{rank}(\chi)$, through which π passes, is greater than both $\mathsf{rank}(\psi)$ and $\mathsf{rank}(\psi')$. It follows that in order to prove that $\mathcal{A}_\varphi$ has countdown depth k it suffices to decompose X into disjoint $X_1, ..., X_k$ such that no X_i contains a pair of $\preceq$-comparable formulae. Define:

$$X_i = \{\psi \in X \mid \text{nesting of countdown operators in } \psi \text{ equals } i\}$$

for every $1 \leq i \leq k$. This exhausts the entire X: each $\psi \in X$ begins with a countdown operators (so the nesting is at least one) and φ (and hence also its subformulae) have nesting at most k. Clearly no two formulae in X_i can be $\preceq$-comparable, so this proves (1) $\implies$ (2).

(2) $\implies$ (1)
Given a countdown automaton $\mathcal{A}$ of countdown depth k, we construct an equivalent formula of $\mu^{<\infty}$-ML with nesting of the countdown operators at most k. The translation $\mathcal{A} \mapsto \varphi_{\mathcal{A}}$ presented in Subsection 4.4.2 is not sufficient: the output formula $\varphi_{\mathcal{A}}$ is equivalent to $\mathcal{A}$ but can have nesting greater than the countdown depth of $\mathcal{A}$. We illustrate this issue with the following toy example.

Example 4.7.9. Consider the following automaton $\mathcal{A}$:

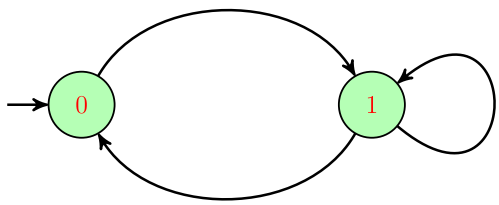

with two states 0 and 1, and 0 being the initial one. The automaton only has ε-transitions. Both states are equal to their ranks $\mathcal{R} = \mathcal{D} = \{0 < 1\}$, which are nonstandard with some fixed initial counter value $\alpha \in \mathsf{Ord}$. Although the language of the automaton is empty and can be described with the trivial formula $\bot$, the translation from Subsection 4.4.2 gives a formula $\varphi_{\mathcal{A}}$:

$$\varphi_{\mathcal{A}} = \mu^{\alpha} x_0.\mu^{\alpha} x_1.\mu^{\alpha} x_0.x_1.$$

The nesting of the μ^{α} operators in $\varphi_{\mathcal{A}}$ equals 3 and is therefore greater than the countdown depth of $\mathcal{A}$ (which is 2).

The reason why the formula $\varphi_{\mathcal{A}}$ has nesting greater than the countdown depth of $\mathcal{A}$ is the use of substitutions. We now introduce a refined version $\mathcal{A} \mapsto \varphi'_{\mathcal{A}}$ of the translation $\mathcal{A} \mapsto \varphi_{\mathcal{A}}$ that avoids that issue. Assume the same setting and notation as in the proof of Theorem 4.4.5. Additionally, assume that $\mathcal{A}$ has countdown depth k witnessed by the decomposition $\mathcal{D} = \mathcal{D}_1 \sqcup \ldots \sqcup \mathcal{D}_k$ of its nonstandard ranks. The original construction proceeds by induction on the ranks of $\mathcal{A}$: for each $r \in \mathcal{R}$ and $q \in Q$ we have an appropriate formula $\psi_{r,q}$ so that at the end $\psi_{r_{\mathsf{max}},q_I} = \varphi_{\mathcal{A}}$ is equivalent to $\mathcal{A}$. Let us construct an enhanced $\psi'_{r,q}$ as follows.

- The base case is not changed:

  $$\psi'_{0,q} = \psi_{0,q}$$

  for all $q \in Q$.

- In the inductive step for rank $r+1$ denote the states with rank r by Q_r. For $q \in Q_r$, instead of:

  $$\psi_{r+1,q} = \eta^{\alpha}_{q}(x_p)_{p \in Q_r}.(\psi_{r,p})_{p \in Q_r}$$

  we take:

  $$\psi'_{r+1,q} = \eta^{\alpha}_{q}(x_p)_{p \in Q_{r,q}}.(\psi'_{r,p})_{p \in Q_{r,q}}$$

  where $Q_{r,q} \subseteq Q_r$ is the least set such that (i) $q \in Q_{r,q}$ and (ii) for every $p, p' \in Q_r$, if $p \in Q_{r,q}$ and $x_{p'} \in \mathsf{FreeVar}(\psi'_{r,p})$ then $p' \in Q_{r,q}$. The intuition is that we only take the coordinates which are reachable from the initial q. For $q \notin Q_r$, instead of:

$$\psi_{r+1,q} = \psi_{r,q}[x_{p_1} \mapsto \psi_{r+1,p_1}, ..., x_{p_d} \mapsto \psi_{r+1,p_d}]$$

we take:

$$\psi'_{r+1,q} = \mu^{\infty}_{y} \langle y, x_{s_1}, ..., x_{s_l} \rangle . \langle \psi'_{r,q}, \psi'_{r+1,s_1}, ..., \psi'_{r+1,s_d} \rangle$$

where $x_{s_1}, ..., x_{s_l} \subseteq x_{q_1}, ..., x_{q_d}$ are those variables x_{q_i} which actually appear in $\psi'_{r,q}$, meaning that $x_{q_i} \in \mathsf{FreeVar}(\psi'_{r,q})$, and y is a fresh variable not used anywhere else.

We put $\varphi'_{\mathcal{A}} = \psi'_{r_{\mathsf{max}},q_I}$, analogously to the original construction.

Let us prove that the enhanced construction does the job: it produces a formula equivalent to the automaton but with low nesting of the countdown operators. Denote:

$$Y = \{\psi'_{r+1,q} \mid \mathsf{rank}(q) = r\}.$$

It is an invariant of the refined construction that for every $r \in \mathcal{R}$ and $q \in Q$, the formula $\xi = \psi'_{r,q}$ has the following properties.

1. For every $\theta \prec \xi$ with either $\theta = \psi'_{r',q'}$ or $\theta = x_{q'} \in \mathsf{FreeVar}(\xi)$ there is a path $q \xrightarrow{\pi} q'$ in $\mathcal{A}$ not visiting ranks greater than:

$$\max\{\mathsf{rank}(q), \mathsf{rank}(q'), r-1\}.$$

2. For every $\theta \prec \xi$ with $\theta = \psi'_{r',q'}$:

$$r' \leq r$$

and the inequality is strict whenever $\xi \in Y$.

3. The formulae $\psi'_{r,q}$ and $\psi_{r,q}$ are equivalent.

In the above, all the equalities between formulae are understood up to isomorphism.

The invariant is clearly true for the base case $r = 0$ so consider the inductive step for $r + 1$.

The Case with $q \in Q_r$. Let us first inspect the situation when $q \in Q_r$ and thus:

$$\xi = \psi'_{r+1,q} = \eta^{\alpha}_{q}(x_p)_{p \in Q_{r,q}}.(\psi'_{r,p})_{p \in Q_{r,q}}.$$

Consider $\theta \prec \xi$ and hence $\theta \preceq \psi'_{r,p}$ for some $p \in Q_{r,q}$.

Towards (1) for ξ assume that either $\theta = \psi'_{r',q'}$ or $\theta = x_{q'} \in \mathsf{FreeVar}(\xi)$. By definition of $Q_{r,q}$ there are $\psi'_{r,p_1}, ..., \psi'_{r,p_k}$ with $p_1 = q, p_k = p$ and $x_{p_{i+1}} \in$

$\mathsf{FreeVar}(\psi'_{r,p_i})$ for each $i < k$. Hence, (1) applied to each ψ'_{r,p_i} gives us paths in $\mathcal{A}$:

$$p_1 \xrightarrow{\pi_1} p_2 \xrightarrow{\pi_2} \dots \xrightarrow{\pi_{k-2}} p_{k-1} \xrightarrow{\pi_{k-1}} p_k,$$

each without ranks greater than $\max\{\mathsf{rank}(p_i), \mathsf{rank}(p_{i+1}), r-1\} = r$. If $\theta = \psi'_{r,p}$ then the composition $\pi_1 \dots \pi_{k-1}$ leading from q to q' witnesses (1) for ξ and θ. Otherwise $\theta \prec \psi'_{r,p}$ and again by (1) we get a path $p \xrightarrow{\pi'} q'$ without ranks greater than $\max\{\mathsf{rank}(p), \mathsf{rank}(q'), r-1\} = \max\{r, \mathsf{rank}(q')\}$. Together the paths compose into $\pi_1 \dots \pi_{k-1}\pi'$ from q to q', therefore witnessing (1) for ξ and θ.

For (2) assume $\theta = \psi_{r',q'}$. Since $\xi \in Y$ we need to prove $r' < r + 1$. If $\theta = \psi'_{r,p}$ then $r = r'$. Otherwise $\theta \prec \psi'_{r,p}$ and $r' \leq r$ by the induction hypothesis.

For (3) observe that $\psi''_{r+1,q} = \eta^{\alpha}_{q}(x_p)_{p \in Q_r}.(\psi'_{r,p})_{p \in Q_r}$ (i.e. the same as $\psi'_{r+1,q}$ except that with full Q_r and not $Q_{r,q}$) is equivalent to $\psi_{r+1,q}$ thanks to (3) applied to all $\psi'_{r,p}$. Whenever $\psi'_{r,p}$ is reachable from the root in $\psi''_{r+1,q}$, then $p \in Q_{r,q}$. It follows that the reachable parts of $\psi''_{r+1,q}$ and $\psi'_{r+1,q}$ are isomorphic and so the formulae are equivalent.

The Case with $q \notin Q_r$. Let us have a look at the other case $q \notin Q_r$ which implies:

$$\xi = \psi'_{r+1,q} = \mu^{\infty}_{y}\langle y, x_{s_1}, \dots, x_{s_l}\rangle.\langle \psi'_{r,q}, \psi'_{r+1,s_1}, \dots, \psi'_{r+1,s_d}\rangle$$

with $x_{s_1}, \dots, x_{s_l} = \mathsf{Var}_r \cap \mathsf{FreeVar}(\psi'_{r,q})$. Assume $\theta \prec \xi$. Similarly to the previous case, either $\theta \preceq \psi'_{r,q}$ or $\theta \preceq \psi'_{r+1,p}$ for some $x_p \in \mathsf{Var}_r \cap \mathsf{FreeVar}(\psi'_{r,q})$.

To prove (1) for ξ assume $\theta = \psi'_{r'q'}$ or $\theta = x_{q'} \in \mathsf{FreeVar}(\xi)$. We need a path from q to q' without ranks greater than $\max\{\mathsf{rank}(q), \mathsf{rank}(q'), r\}$. If $\theta \preceq \psi'_{r,q}$ then either $\theta = \psi_{r,q}$ and we are done with the trivial path $\pi = q$ or $\theta \prec \psi'_{r,q}$ and the induction hypothesis for $\psi'_{r,q}$ gives us the desired $q \xrightarrow{\pi} q'$. The remaining case is when $\theta \preceq \psi'_{r+1,p}$ for $x_p \in \mathsf{Var}_r \cap \mathsf{FreeVar}(\psi'_{r,q})$. Since x_p is free in $\psi'_{r,q}$, (1) gives us $q \xrightarrow{\pi} p$ with ranks not greater than $\max\{\mathsf{rank}(q), \mathsf{rank}(p), r-1\}$. If $\theta = \psi'_{r+1,p}$ and consequently $q' = p$ we are done. Otherwise $\theta \prec \psi'_{r+1,p}$ and again by (1) we get $p \xrightarrow{\pi'} q'$ with ranks at most $\max\{\mathsf{rank}(p), \mathsf{rank}(q'), r\}$. Since $\mathsf{rank}(p) = r$ it follows that the composition $q \xrightarrow{\pi} p \xrightarrow{\pi'} q'$ has ranks at most $\max\{\mathsf{rank}(q), \mathsf{rank}(q'), r\}$, as desired.

For (2) assume that $\theta = \psi'_{r'q'}$. Since $\xi \notin Y$, we only need $r' \leq r + 1$. Either $\theta \preceq \psi'_{r,q}$ or $\theta \preceq \psi'_{r+1,p}$. If $\theta = \psi'_{r,q}$ or $\theta = \psi'_{r+1,p}$ we are done. Otherwise θ is a proper subformula $\theta \prec \psi'_{r,q}$ or $\theta \prec \psi'_{r+1,p}$ and we get the claim from the induction hypothesis or the previous case (i.e. with $q \in Q_r$), respectively.

Towards (3) consider $\psi''_{r+1,q} = \psi_{r,q}[x_{p_1} \mapsto \psi'_{r+1,p_1}, \dots, x_{p_d} \mapsto \psi'_{r+1,p_d}]$. By the induction hypothesis, $\psi''_{r+1,q}$ and $\psi_{r+1,q}$ are equivalent. The equivalence of $\psi''_{r+1,q}$ and ξ follows by a straightforward analysis of the induced semantic games (a formalist reader could invoke Proposition 2.2.2). This completes the proof of the invariant (1), (2) and (3).

To complete the proof of the implication (2) $\implies$ (1) we show that $\varphi'_{\mathcal{A}}$ has nesting of the countdown operators at most k. Assume towards contradiction that there are:

$$\theta_1 \prec \theta_2 \prec ... \prec \theta_{k+1} \preceq \varphi'$$

with each θ_i beginning with a countdown operator. Every subformula θ_i of $\varphi'_{\mathcal{A}}$ beginning with a countdown operator must be isomorphic to $\psi'_{r_i+1,q_i} \in Y$ with nonstandard $r_i = \mathsf{rank}(q_i)$. Thus, for every $i \leq k$ (2) implies $r_i < r_{i+1}$. By the pigeonhole principle some $r_i < r_{i'}$ belong to the same $\mathcal{D}_n$. By (1), there is a path in $\mathcal{A}$ from q_i to $q_{i'}$ not visiting ranks greater than r_i and $r_{i'}$, which is a contradiction.

(2) $\implies$ (3)
In the proof of Theorem 4.7.4 we translated a countdown automaton $\mathcal{A}$ with nonstandard ranks $\mathcal{D} = \{r_1, ..., r_l\}$ to an automaton with stacked counters $\mathcal{A}'$ storing the counter values $\mathsf{ctr}(r_1), ..., \mathsf{ctr}(r_l)$ of any counter assignment $\mathsf{ctr} : \mathcal{D} \to \mathsf{Ord}$ on its stack. The height of the stack of $\mathcal{A}'$ was therefore bounded by $|\mathcal{D}|$. However, instead of keeping all the counter values on the stack it suffices if $\mathcal{A}'$ only stores the counters with non-initial values. If $\mathcal{A}$ has countdown depth k, then in any its accessible configuration at most k counters have non-initial values. As a consequence, such more succinct representation of the counter assignments of $\mathcal{A}$ requires a stack of height at most k.

(3) $\implies$ (2)
This case is even easier than the previous one. The translation from automaton $\mathcal{A}$ with stacked counters to countdown automaton $\mathcal{A}'$ given in the proof of Theorem 4.7.4 does not require any modification: if $\mathcal{A}$ has a stack bounded by k then the constructed $\mathcal{A}'$ has countdown depth at most k. This is witnessed by the decomposition of the nonstandard ranks $\mathcal{D} = \{t^i_\exists, t^i_\forall \mid 1 \leq i \leq k\}$ of $\mathcal{A}'$ into disjoint $\mathcal{D}_1, ..., \mathcal{D}_k$ with:

$$\mathcal{D}_i = \{t^i_\exists, t^i_\forall\}$$

for all $1 \leq i \leq k$. Any nonstandard rank t^i_P in $\mathcal{A}'$ can only be visited if the current state of the automaton remembers a stack shape of height i with P on top. Thus, every path between $t^i_\exists$ and $t^i_\forall$ must pass through an ε-transition $p \xrightarrow{\varepsilon} q$ in $\mathcal{A}'$ corresponding to a pop-transition $p \xrightarrow{\mathsf{pop}} q$ in $\mathcal{A}$ poping the stack from height i to $i-1$. The state q is not a source of a stack-transition in $\mathcal{A}$ (for it is a target of such a transition and $\mathcal{A}$ does not fire two stack-transitions in a row) and so $r = \mathsf{rank}'(q) = \mathsf{rank}(q)$. Since $q \in Q^{i-1}$, the rank $r \in \mathcal{R}^{i-1}$ is greater than both $t^i_\exists$ and $t^i_\forall$. This completes the proof of Theorem 4.7.8. $\square$

4.8 Strictness of the Countdown Complexity Hierarchy

The three formalisms: countdown logic, countdown automata and automata with stacked counters, are equivalent. In the previous section we introduced a measure of complexity for each: nesting of the countdown operators, countdown depth and maximal height of the stack, respectively. Theorem 4.7.8 establishes equivalence of the measures: restricting any of the formalisms by only considering its instances of complexity k results in the same expressive power as the restriction of any other with the same k. This induces a natural stratification of languages into classes. Given a language L we define its *countdown complexity* to be the least $k < \omega$ such that L is defined by a $\mu^{<\infty}$-ML sentence with nesting of the countdown operators k (or ∞ if no sentence defines L).

Such countdown complexity is an arguably robust measure. So far we have not shown, however, that it is substantial in the sense that the hierarchy of languages with greater and greater complexity does not collapse. We now prove that under mild assumptions the hierarchy is strict. For the rest of this section assume that the only ordinal appearing in formulae is ω (we will discuss how this assumption can be weakened in Remark 4.8.4 at the end of the section).

The existence of languages not definable in $\mu^{<\infty}$-ML is immediate: if we restrict our attention to finite models, $\mu^{<\infty}$-ML and μ-ML have the same expressive power and clearly not all languages of finite models are definable in the later. The following theorem expresses the non-trivial content of the mentioned strictness.

Theorem 4.8.1. *For every $k < \omega$, there are languages of countdown complexity k.*

Proof. For all k, we provide examples of languages definable with countdown nesting $k+1$ but not k. In order to prove strictness, it suffices to prove it on a restricted class of models. From now on, focus on the monomodal case with no colors (i.e. $|\mathsf{Act}| = 1$ and $\mathsf{Prop} = \emptyset$). In this setting models consist of universe with a single binary relation. We will show that the hierarchy is strict already on the class of transitive, linear, well-founded models. Up to isomorphism, these are just ordinals and hence we confine our attention to ordinal models, as defined in Definition 3.2.1 (with the special case $\mathsf{Prop} = \emptyset$).

Since κ is an induced submodel of κ' whenever $\kappa \leq \kappa'$, we can consider a single ordinal model with κ big enough. For our purposes, the first uncountable ordinal ω_1 is sufficient.

We call a subset $S \subseteq \omega_1$ *stable* on an interval $I \subseteq \omega_1$ if either $S \cap I = I$ or $S \cap I = \emptyset$. S is *stable above* α if it is stable on the interval $[\alpha, \omega_1)$. A *stabilization point* of a valuation $\mathsf{val} : \mathsf{Var} \to \mathcal{P}(\omega_1)$ is the least $\alpha \leq \omega_1$ such that interpretations of all the variables are stable above α.

Observe that the set $[\omega^k, \omega_1) \subseteq [0, \omega_1)$ can be defined by the following sentence with countdown nesting k:

$$[\omega^k, \omega_1) = [\![\nu^\omega x_1 ... \nu^\omega x_k . \Diamond(\bigwedge_{i \le k} x_i)]\!]. \tag{4.25}$$

Indeed, the semantic game can be decomposed into two alternating steps: (i) $\forall$dam chooses a tuple of finite ordinals $(\alpha_1, ..., \alpha_k) \in \omega^k$ and (ii) $\exists$ve responds with a successor in the model. Since at each step $\forall$dam has to pick a lexicographically smaller tuple (and he starts by picking any tuple) it is easy to see that he wins iff the initial point is at least ω^k. We will show that for all $k > 0$, countdown nesting k is *necessary* to define this language. The proof relies on the following lemma.

Lemma 4.8.2. *For every $k < \omega$ and a formula φ with countdown nesting k, there exists an ordinal $\alpha_\varphi < \omega^{k+1}$ such that φ stabilizes α_φ above the valuation, meaning that for every valuation* val *stabilizing at β, $[\![\varphi]\!]^{\mathsf{val}}$ is stable above $\beta + \alpha_\varphi$.*

From this the theorem follows immediately, as the sentence φ has no free variables and thus it stabilizes at $\alpha_\varphi < \omega^{k+1}$ regardless of the valuation. Hence, let us prove the lemma.

Proof. We start with the following proposition.

Proposition 4.8.3. *For every countdown formula φ there is a finite constant $t_\varphi < \omega$ such that for every valuation* val *stable above κ, in the part $[\kappa, \omega_1)$ of the model above κ, φ changes its truth value at most t_φ times.*

Proof. Since without loss of generality the formula is guarded (see Proposition 4.5.3), by Proposition 4.5.4 we may assume that in the semantic game $\exists$ve always uses a pre-modally counter-independent strategy. But the number z of possible pre-modal components for such strategies is finite, so if φ changed its value more than $t_\varphi = 2z + 1$ times above κ, there would be $\kappa \le \alpha < \zeta < \beta$ such that $\exists$ve wins from α and β with the same pre-modal component, but loses from ζ in between, which is impossible by Proposition 4.5.6. $\square$

We prove Lemma 4.8.2 by induction on the complexity of the formula φ. The base case is immediate, as for every $x \in \mathsf{Var}$ it suffices to take $\alpha_x = 0$. For propositional connectives and modal operators we take $\alpha_{\psi_1 \vee \psi_2} = \alpha_{\psi_1 \wedge \psi_2} = \mathsf{max}(\alpha_{\psi_1}, \alpha_{\psi_2})$ and $\alpha_{\Diamond\psi} = \alpha_{\Box\psi} = \alpha_\psi + 1$. The remaining non-trivial cases are countdown and fixpoint operators.

- Assume $\varphi = \eta_i^\omega \overline{x}.\overline{\psi}$ with $\overline{x} = \langle x_1, ..., x_d \rangle$ and $\overline{\psi} = \langle \psi_1, ..., \psi_d \rangle$. Let $\Phi = \{\theta_1, ..., \theta_l\}$ be the set of all maximal subformulae of $\overline{\psi}$ not using any variable from $\overline{x}$. For each θ, pick a fresh variable y_θ and put:

$$\psi'_j = \psi_j[\theta_1 \mapsto y_{\theta_1} ... \theta_l \mapsto y_{\theta_l}]$$

i.e. starting from the root ψ_j, we replace every subformula θ that has no variables from $\overline{x}$ with a fresh variable y_θ.[2] Observe that $\psi_j = \psi'_j[y_{\theta_1} \mapsto \theta_1 ... y_{\theta_l} \mapsto \theta_l]$, so:

$$\eta_i^\omega \overline{x}.\overline{\psi} \equiv (\eta_i^\omega \overline{x}.\overline{\psi'})[y_{\theta_1} \mapsto \theta_1 ... y_{\theta_l} \mapsto \theta_l].$$

Note that if φ has countdown nesting at most k then each ψ'_j and each θ has countdown nesting less than k. Thus, by the induction hypothesis there exist $\alpha_{\psi'_j} < \omega^k$ and $\alpha_\theta < \omega^k$ s.t. ψ'_j and θ stabilize $\alpha_{\psi'_j}$ and α_θ above the valuation, respectively. Denote $\alpha_{\overline{\psi'}} = \mathsf{max}\{\alpha_{\psi'_1}, ..., \alpha_{\psi'_d}\}$.

For $m < \omega$, consider the m-th unfolding given by $\psi_j'^0 = x_j$ and $\psi_j'^{m+1} = \psi'_j[x_1 \mapsto \psi_1'^m ... x_d \mapsto \psi_d'^m]$. It follows by a straightforward induction on m that each $\psi_j'^m$ is stable $\alpha_{\overline{\psi'}} \times m$ above the valuation. Moreover, for any valuation val we have:

$$[\![\mu_j^\omega \overline{x}.\overline{\psi'}]\!]^{\mathsf{val}} = \bigcup_{m<\omega} [\![\psi_j'^m]\!]^{\mathsf{val}} \qquad \text{and} \qquad [\![\nu_j^\omega \overline{x}.\overline{\psi'}]\!]^{\mathsf{val}} = \bigcap_{m<\omega} [\![\psi_j'^m]\!]^{\mathsf{val}}$$

so $\eta_j^\omega \overline{x}.\overline{\psi'}$ is stable $\alpha_{\overline{\psi'}} \times \omega$ above the valuation. Finally, we obtain that $\varphi = (\eta_j^\omega \overline{x}.\overline{\psi'})[\theta_1 \mapsto y_{\theta_1} ... \theta_l \mapsto y_{\theta_l}]$ is stable α_φ above valuation with:

$$\alpha_\varphi = \mathsf{max}\{\alpha_{\theta_1}, ..., \alpha_{\theta_l}\} + \alpha_{\overline{\psi'}} \times \omega.$$

Since $\alpha_{\overline{\psi'}} \times \omega < \omega^{k+1}$ and for each θ, $\alpha_\theta < \omega^k$ it follows that $\alpha_\varphi < \omega^{k+1}$.

- Assume $\varphi = \eta_i^\infty \overline{x}.\overline{\psi}$ with $\overline{x} = \langle x_1, ..., x_d \rangle$ and $\overline{\psi} = \langle \psi_1, ..., \psi_d \rangle$. Note that we give a proof for the general, vectorial case. Although η^∞ is a fixpoint operator and so it could be replaced with scalar ones using the Bekić principle (2.9), such rewriting could increase the nesting of the countdown operators.

  The countdown nesting of each ψ_j is not greater than that of φ. For each $j \leq d$, let $t_{\eta_j^\infty \overline{x}.\overline{\psi}} < \omega$ be the constant from Proposition 4.8.3 and $\alpha_{\psi_j} < \omega^{k+1}$ the constant that exists by the inductive hypothesis. Put $\alpha_{\overline{\psi}} = \mathsf{max}_{j \leq d}(\alpha_{\psi_j})$, $t_{\mathsf{max}} = \mathsf{max}_{j \leq d}(t_{\eta_j^\infty \overline{x}.\overline{\psi}})$. We define:

$$\alpha_\varphi = \alpha_{\overline{\psi}} \times t_{\mathsf{max}} \times d.$$

  Clearly $\alpha_\varphi < \omega^{k+1}$, as $\alpha_{\overline{\psi}} < \omega^{k+1}$ and $t_{\mathsf{max}} < \omega$. It therefore suffices to show that $[\![\varphi]\!]^{\mathsf{val}}$ stabilizes above α_φ above val. Define the valuation:

$$\mathsf{val}_\varphi = \mathsf{val}[x_1 \mapsto \eta_1^\infty \overline{x}.\overline{\psi}^{\mathsf{val}} \; ... \; x_d \mapsto \eta_d^\infty \overline{x}.\overline{\psi}^{\mathsf{val}}]$$

[2] Recall that we do not identify isomorphic subformulae, and so there are no substitutions *inside* the θ's. In particular, the order of substitutions does not matter.

and let κ be the stabilization point of val. It suffices to prove that val_φ is stable above $\kappa + \alpha_\varphi$. For y not in $\overline{x}$, $\mathsf{val}(y)$ and thus also $\mathsf{val}_\varphi(y)$ is already stable above κ. For $j \leq d$, $\mathsf{val}_\varphi(x_j)$ changes value above κ at most t_{max} times. Together this means that val_φ changes its value at most $t_{\mathsf{max}} \times d$ times above κ.

On the other hand, if val_φ does not change its value for at least $\alpha_{\overline{\psi}}$ steps, it remains stable forever, i.e. if for some $\kappa \leq \alpha < \omega_1$ we have that val_φ is stable on the interval $[\alpha, \alpha + \alpha_{\overline{\psi}}]$, then it is stable on the entire $[\alpha, \omega_1)$ (i.e. stable above α). Once we prove this, correctness of α_φ follows: after at most $t_{\mathsf{max}} \times d$ blocks, each of length at most $\alpha_{\overline{\psi}}$, the valuation val_φ stabilizes.

Assume that val_φ is stable on $[\alpha, \alpha + \alpha_{\overline{\psi}}]$. We prove:

$$\mathsf{val}_\varphi \text{ is stable on } [\alpha, \beta) \implies \mathsf{val}_\varphi \text{ is stable on } [\alpha, \beta] \tag{4.26}$$

for all $\beta > \alpha + \alpha_{\overline{\psi}}$. Given the above, we get by induction on $\beta > \alpha + \alpha_{\overline{\psi}}$ that val_φ is stable on $[\alpha, \beta)$.

Since $\alpha_{\psi_j} \leq \alpha_{\overline{\psi}}$, the semantics $[\![\psi_j]\!]^{\mathsf{val}_\varphi}$ is stable $\alpha_{\overline{\psi}}$ over val_φ. Since modal formulae do not look backwards, we can strengthen this:

$$\mathsf{val}_\varphi \text{ is stable on } [\alpha, \beta] \implies [\![\psi_j]\!]^{\mathsf{val}_\varphi} \text{ is stable on } [\alpha + \alpha_{\overline{\psi}}, \beta].$$

for every $\beta > \alpha + \alpha_{\overline{\psi}}$. Moreover, since all $\overline{x}$ are guarded in ψ_j and for $y \notin \overline{x}$ the valuation $\mathsf{val}_\varphi(y)$ is stable already above $\kappa \leq \alpha$, we can further strengthen the implication by weakening the premise:

$$\mathsf{val}_\varphi \text{ is stable on } [\alpha, \beta) \implies [\![\psi_j]\!]^{\mathsf{val}_\varphi} \text{ is stable on } [\alpha + \alpha_{\overline{\psi}}, \beta].$$

for every $\beta > \alpha + \alpha_{\overline{\psi}}$. For each $j \leq d$ we have:

$$\begin{aligned} \mathsf{val}_\varphi(x_j) &= [\![\eta_j^\infty \overline{x}.\overline{\psi}]\!]^{\mathsf{val}} \\ &= [\![\eta_j^\infty \overline{x}.\overline{\psi}]\!]^{\mathsf{val}_\varphi} \\ &= [\![\psi_j]\!]^{\mathsf{val}_\varphi} \end{aligned}$$

The first equality is the definition of val_φ. The second one follows from the observation that the valuation for $\overline{x}$ is irrelevant for the semantics of $\eta_j^\infty \overline{x}.\overline{\psi}$. The third one is due to η^∞ being a fixpoint operator.

Combining the equality and the last implication we get:

$$\mathsf{val}_\varphi(x_j) \text{ is stable on } [\alpha, \beta) \implies \mathsf{val}_\varphi(x_j) \text{ is stable on } [\alpha + \alpha_{\overline{\psi}}, \beta].$$

for all $\beta > \alpha + \alpha_{\overline{\psi}}$. Since $[\alpha, \beta)$ and $[\alpha + \alpha_{\overline{\psi}}, \beta]$ overlap, stability on both implies stability on their union $[\alpha, \beta]$. Because x_j above is arbitrary and $\mathsf{val}_\varphi(y)$ is already stable above α for $y \notin \overline{x}$ we obtain (4.26).

This finishes the proof of Lemma 4.8.2 and Theorem 4.8.1. $\square$

$\square$

Remark 4.8.4. Above we assumed that the only ordinal appearing in formulae is ω. However, the proof works under a weaker assumption: there is a maximal ordinal α used in formulae and this α is *additively indecomposable* (meaning that if $\beta, \beta' < \alpha$ then $\beta + \beta' < \alpha$). Every cardinal number is additively indecomposable. Another example could be ω^k for natural k.

4.9 Decidability Issues

In this section we discuss decidability issues in the countdown μ-calculus. Note that in a finite model every monotone map reaches its fixpoints in finitely many steps. Hence, if we replace every η^α in φ with η^∞ and denote the resulting formula by φ^{st}, then in every *finite* model $[\![\varphi]\!] = [\![\varphi^{\mathsf{st}}]\!]$. It immediately follows that:

Proposition 4.9.1. *The model checking problem for the* $\mu^{<\infty}$-ML*, i.e. the problem: "Given* $\varphi \in \mu^{<\infty}$-$\mathsf{ML}$ *and a point* m *in a (finite) model* $\mathcal{M}$*, does* $\mathsf{m} \models \varphi$*?" is decidable.*

Note that as a corollary we get that deciding the winner of a given (finite) countdown game $\mathcal{G}$ is also decidable, as the set of positions where $\exists$ve wins can be easily defined in $\mu^{<\infty}$-ML.

A more interesting problem is *satisfiability*: "Given $\varphi \in \mu^{<\infty}$-$\mathsf{ML}$, is there a model $\mathcal{M}$ and a point m such that $\mathsf{m} \models \varphi$?". Satisfiability is closely related to *validity*, the question whether a given formula is satisfied in all models. A formula is satisfiable iff its negation is not valid. We conjecture that satisfiability (and therefore also validity) is decidable for the full logic.

Conjecture 4.9.2. *The satisfiability problem for* $\mu^{<\infty}$-ML *is decidable.*

Conjecture 4.9.2 seems challenging. For now, we present some of its special cases.

Proposition 4.9.3. *A formula* $\varphi \in \mu^{<\infty}$-$\mathsf{ML}$ *has* positive countdown *if it does not use* ν^α *with* $\alpha \neq \infty$*. The satisfiability problem is decidable for such formulae.*

Proof. Observe that for φ with positive countdown, in every model we have $[\![\varphi]\!] \subseteq [\![\varphi^{\mathsf{st}}]\!]$. Hence, if φ is satisfiable, then so is φ^{st}. But since μ-ML has a finite model property, this means that φ^{st} has a *finite* model, where φ^{st} and φ are equivalent. Thus, φ is satisfiable iff φ^{st} is, and the problem reduces to μ-ML satisfiability. $\square$

The negation of a formula with positive countdown has *negative countdown*, i.e. $\alpha = \infty$ for every μ^α. Thus, dualizing the above we get that the validity problem is decidable for such formulae.

The finite model property of positive formulae makes the satisfiability problem easier to solve but also less interesting. On the contrary, Example 4.2.4 demonstrates that formulae using ν^α with $\alpha \neq \infty$ lack the finite model property. The next subsection is devoted to solving a fragment of the logic that allows for some negative countdown: the Büchi fragment.

4.9.1 Büchi Countdown Automata over Infinite Words

In the classical setting (as found in [25]), a simple parity game or automaton is called *Büchi* if it has only two ranks $r^\exists < r^\forall$, the most important $r^\forall$ belonging to $\forall$dam and the other $r^\exists$ belonging to $\exists$ve. We extend that definition directly to the countdown setting. We call a countdown game or a countdown automaton Büchi if they have only two such ranks $r^\exists < r^\forall$.

We solve the satisfiability problem for Büchi countdown automata over infinite words (seen as monomodal models the same way as in Example 4.2.5). The case when $r_\forall$ is standard is already covered by Proposition 4.9.3. We therefore only consider the case when $r^\forall$ is nonstandard. Moreover, for simplicity we assume that the initial counter value for $r^\forall$ equals ω. On the other hand, the next remark says that in the case of infinite words it does not matter if $r^\exists$ is standard or not. Hence, for the sake of simpler notation we assume that $r^\exists$ is standard.

Remark 4.9.4. If the models under consideration are finitely branching, so are the arenas of the corresponding semantic games. In such case it does not matter whether $r^\exists$ is standard or not so we assume it is standard.

To see that standardness of $r^\exists$ does not matter we show that if $\exists$ve has a winning strategy σ with standard $r^\exists$ then she may use this σ to win in the harder case with nonstandard $r^\exists$. Without loosing generality, σ is configurational (i.e. dictates the same move after plays leading to the same configuration). This implies that in any σ-reachable configuration γ there is a finite bound $k_\gamma < \omega$ on the lengths of σ-plays from γ that only contain $r^\exists$ but not $r^\forall$. Otherwise, the tree of all such plays would be an infinite but finitely branching tree and so by König's Lemma it would contain an infinite σ-play containing only $r^\exists$, which contradicts that σ is winning from γ.

Using these bounds k_γ, $\exists$ve may win in the harder case with nonstandard $r^\exists$: she makes the same positional moves as originally and whenever in a configuration γ she needs to pick a value for the counter corresponding to $r^\exists$ she chooses k_γ.

Satisfiability of Büchi automata over infinite words is the most technically involved special case for which we managed to prove Conjecture 4.9.2.

Theorem 4.9.5. *The satisfiability problem for countdown Büchi automata over infinite words is decidable.*

Proof. In [8] the author introduce a fragment of the $\mathsf{MSO}+\mathsf{U}$ logic called S-formulae. This is the least subset of $\mathsf{MSO}+\mathsf{U}$ containing all MSO formulae and closed under $\vee, \wedge, \exists, \forall$ and U. The satisfiability problem for S-formulae over infinite words is decidable [9]. Therefore, Theorem 4.9.5 follows from the following proposition.

Proposition 4.9.6. *Every language of infinite words recognized by a Büchi countdown automaton $\mathcal{A}$ is recognized by an S-formula computable from $\mathcal{A}$.*

We will first show some useful properties of the Büchi countdown games. We start with a limited form of positionality. Since there is only one non-standard rank, we simplify the notation and identify the counter assignment $\mathsf{ctr} : \{r^\forall\} \to \mathsf{Ord}$ with the only value $\mathsf{ctr}(r^\forall)$. For every ordinal α, denote the set of all the winning configurations with counter value at most α by Conf_α. The next proposition says that positional strategies are sufficient against finite counter values.

Proposition 4.9.7. *Assume a Büchi countdown game. For every $\alpha < \omega$ there is a positional strategy σ_α for $\exists$ve winning from every configuration in Conf_α.*

Proof. Fix $\alpha < \omega$ and let σ be a configurational strategy winning from every winning configuration. We construct a positional σ_α winning from every configuration in Conf_α. For every position v appearing in some configuration in Conf_α consider all the counter values β such that $(v, \beta, \mathsf{psn}) \in \mathsf{Conf}_\alpha$. No such β is greater than α and so for each v there must be the greatest $\beta_v \leq \alpha$. Then, $\sigma_\alpha(v) = \sigma(v, \beta_v, \mathsf{psn})$, i.e. in v $\exists$ve assumes the worst possible scenario. Such σ_α is well-defined and always dictates legal moves, because all the configurations reachable from Conf_α by a σ_α-play belong to Conf_α. This is because $r^\forall$ is the most important rank and so the counter never increases and in particular never gets bigger than α. Hence, no finite σ_α-play is lost by $\exists$ve.

It remains to prove that she also wins every infinite σ_α-play π. By well-foundedness of ω and since the counter never increases, after some finite prefix π_0 of π the counter does not change at all. This implies that after that prefix the worst possible values $\beta_{v_1}, \beta_{v_2}, ...$ for the consecutive remaining positions $v_1, v_1, ...$ of π can only get bigger but not smaller. But these values are never greater than α, so from some moment they are constant. However, this implies that from some moment π is actually consistent with σ. Since the later is winning, the parity condition is satisfied and so $\exists$ve wins π. $\square$

Given a Büchi countdown game, Proposition 4.9.7 gives $\exists$ve an infinite sequence of positional strategies winning against greater and greater finite counter values. Although the sequence is infinite and so cannot be directly encoded as a finite labelling of the input, we use it to construct a finitary witness for ∃ve's victory.

Definition 4.9.8. Assume a semantic game induced by a Büchi countdown automaton over input word $\mathsf{w} \in \Sigma^\omega$. A *victory witnessing quadruple* consists of a pair of positional strategies σ_ω and σ_∞, a sequence $l_0 < l_1 < \ldots = \bar{l}$ of indices in w and a subset $W \subseteq V$ of positions such that:

1. If $v \in W$ then either: $v \in V_\forall$ and all its successor positions $E(v)$ are in W; or $v \in V_\exists$, $\sigma_\infty(v)$ is defined and belongs to W.

2. For every maximal σ_ω-play starting from a position in W with index l_i: before reaching index l_{i+1} the play either reaches ∀dam's deadlock or visits $r^\forall$ more than i times.

Such quadruples witnesses that ∃ve wins from W against all finite counter values.

Proposition 4.9.9. *Let $\sigma_\omega, \sigma_\infty, \bar{l}, W$ be a victory witnessing quadruple in $\mathcal{G}(\mathcal{A})$ for w. For every $v \in W$ and counter value $\alpha < \omega$, ∃ve wins from $(v, \alpha, \mathsf{cdn})$.*

Proof. Let i be the least number such that the index of v is not greater than l_i. ∃ve wins as follows. First, until reaching the index $j = \max(k, \alpha)$, play according to σ_∞. The first item (1) guarantees that this way ∃ve does not loose and all the visited positions belong to W. Then, upon reaching index j, switch to σ_ω. Thanks to the second item (2), the rest of the play is guaranteed either to pass through $r^\forall$ more than $j \geq \alpha$ times or to reach ∀dam's deadlock. In either case, ∃ve wins. □

The key fact is that the existence of an appropriate quadruple is not only sufficient but also necessary for ∃ve to win from a given v against all $\alpha < \omega$. There exists a universal witness for all such positions.

Proposition 4.9.10. *Assume a game $\mathcal{G}(\mathcal{A})$ for w as before. There exists a victory witnessing quadruple $\sigma_\omega, \sigma_\infty, \bar{l}, W$ such that $v \in W$ iff ∃ve wins from $(v, \alpha, \mathsf{cdn})$ for every $\alpha < \omega$.*

Proof. Given W from which ∃ve wins against all finite counter values, we construct appropriate $\sigma_\omega, \sigma_\infty$ and $\bar{l}$. Let $S = (\sigma_\alpha)_{\alpha<\omega}$ be the infinite sequence of better and better positional strategies given by Proposition 4.9.7, with σ_α winning against counter value α (and thus also against all the smaller counter values) whenever possible. Observe that if we take any infinite subsequence S' of S then it also does the job: for every counter value α it contains a strategy σ_β with $\beta \geq \alpha$.

A positional strategy for ∃ve is a partial map $\sigma : V_\exists \to V \cup \{\mathsf{undefined}\}$ from positions controlled by ∃ve to positions, subject to some conditions.

If the arena is finitely-branching we can always take a subsequence S' of S convergent to a *limit strategy* σ_∞. In abstract terms, this means that we view the set of all strategies as a topological space (with the standard, product topology on functions) and derive its compactness from finite branching of the arena. Explicitly, convergence means that for every finite $V_0 \subseteq V$, all but finitely many strategies in the sequence agree with σ_∞ on V_0. Be warned that although the limit strategy is well-defined, it does not have to be winning: infinite plays may be lost. Nonetheless, the set W together with σ_∞ satisfy the first condition (1) from the definition.

The other strategy σ_ω and sequence $\bar{l}$ are defined inductively. In the i-th step we want to have bounds $l_0 < ... < l_i$ and the part of σ_ω on positions before l_i (i.e. positions belonging to $[0, l_i) \times Q \subseteq V$) so that (2) is true for all $j < i$. We start with $l_0 = 0$ and σ_ω undefined everywhere. Assume bounds $l_0 < ... < l_i$ and the part of σ_ω on positions before l_i. We define l_{i+1} and the behavior of σ_ω on $[l_i, l_{i+1})$. Consider all σ_i-plays starting in a position in W with index l_i. Every such play π is won if the counter starts with value at most i. This means that π ends with either $i+1$-st visit to $r^\forall$ or $\forall$dam's deadlock. Moreover, there must be a uniform bound $l < \omega$ on the length of all such plays. Without such a bound, by the pigeonhole principle there would be a single position v in which arbitrarily long such plays originate. This would imply that the tree of all σ_i-plays starting at v is infinite and finitely branching (the latter follows from finite branching of the arena) and so by König's Lemma contains an infinite σ_i-play starting at v. This is not possible because $\exists$ve looses all such plays and the strategy σ_i wins against counter value i. Hence, putting $l_{i+1} = l_i + l$ and defining σ_ω on $[l_i, l_{i+1})$ to be equal to σ_i, we get our goal. $\square$

The phase of a Büchi countdown game before the first visit to $r^\forall$ is a simple reachability game: the counter has value ω and $\exists$ve wins iff she manages to arrive at a position of rank $r^\forall$ from which she wins against all finite counter values. It follows from Propositions 4.9.9 and 4.9.10 that she wins from v_I (against the initial counter value ω) iff there exist $\sigma_I, \sigma_\omega, \sigma_\infty, \bar{l}$ and W such that (i) $\sigma_\omega, \sigma_\infty, \bar{l}, W$ is a victory witnessing quadruple and (ii) σ_I is a positional strategy in the simple reachability game obtained from the original $\mathcal{G}(\mathcal{A})$ by replacing $W \cap \mathsf{rank}^{-1}(r^\forall)$ and $(V - W) \cap \mathsf{rank}^{-1}(r^\forall)$ with positions where $\exists$ve immediately wins and looses, respectively.

Positional strategies $\sigma_I, \sigma_\omega, \sigma_\infty$, sets of positions W and indices $\bar{l}$ can be encoded in a straightforward way as subsets of the universe (the number of sets required to encode each object may be greater than one but depends only on $\mathcal{A}$ and is fixed). Moreover the second condition (2) from the definition of a victory witnessing quadruple is equivalent to saying that: for every sequence $\pi_0, \pi_1, ...$ of σ_ω-plays, with each π_i starting from W with index l_i, maximal within interval $[l_i, l_{i+1})$ and not ending in $\forall$dam's deadlock, the numbers of visits to $r^\forall$ in each π_i are unbounded. Such sequence $\overline{\pi}$ can be encoded as a single positional strategy for ∀dam. Hence, the fact that ∃ve wins the game can be described by an S-formula of shape:

$$\exists_{X_1} ... \exists_{X_m} . \forall_{Y_1} ... \forall_{Y_n} \mathsf{U}_Z . \varphi$$

with φ being a plain MSO formula (with no U). This completes the proof of Proposition 4.9.6 and therefore also Theorem 4.9.5. $\square$

Chapter 5

Conclusions

We investigated logics invariant under bisimulation.

Bisimulation-invariant Model Theory. We looked at the classical model-theoretic questions for theories expressed in modal logic. We characterized bisimulational categoricity of modal theories over various classes of models. All the obtained characterizations are simple and essentially equate bisimulational categoricity with existence of an image-finite model for the theory. The proofs, however, require different arguments and the presented counterexamples show that the results cannot be easily generalized to a metatheorem. This suggests several directions for further research.

1. Find a list of reasonable conditions and prove a metatheorem capturing bisimulational categoricity over classes of models satisfying these conditions. A good starting point could be to look at classes that are FO-axiomatizable. Although our counterexamples show that this condition is not sufficient, it helps a lot because it guarantees the existence of models that realize limit behaviors. The key difficulty is then to construct models omitting such a limit behavior.

2. Find characterizations for chosen classes of interest. This is especially challenging for classes that are not FO-axiomatizable and consequently ML need not be compact, which in turn makes the very existence of limit behaviors nontrivial. Nonetheless, the example of ordinal models $\mathcal{C}_{\mathsf{Ord}}$ shows that in some cases difficulties can be overcome.

3. Investigate bisimulational categoricity for bisimulation-invariant logics other than ML, for instance μ-ML. Since such logics are rarely compact even over arbitrary models, this leads to difficulties similar to the ones mentioned in the previous item. Again, this is challenging but not completely hopeless. For instance, one can use a game-theoretic analysis to show that over $\mathcal{C}_{\mathsf{Ord}}$ logics ML and μ-ML are equivalent. Hence, over this class bisimulational categoricity for μ-ML is the same as for the already tamed ML.

The above questions concern bisimulational categoricity, the topic we investigated the most in the model-theoretic part of this dissertation. However, bisimulation-invariant model theory is not limited to bisimulational categoricity. Results such as compactness of ML (with finite signature) over $\mathcal{C}_{\mathsf{Ord}}$ show that there are many other areas awaiting exploration.

Countdown μ-calculus. We introduced and investigated the countdown μ-calculus $\mu^{<\infty}$-ML. We extended the classical (effective) correspondence between logic games and automata by introducing countdown games and automata together with appropriate translations. The connection helped us to establish several facts. Some of these generalize the classical results. For example, formulae of $\mu^{<\infty}$-ML can be always rewritten to a guarded form, and greater nesting of the countdown operators leads to a greater expressive power.

On the other hand, some things differ from the classical case. The vectorial variant of the calculus is more expressive than its scalar fragment. Moreover, our new games are not positional (or even memory-finite) and since the definable languages are provably not closed under projections no simple *nondeterministic* automata can match the logic. Because of this, we cannot simply use standard techniques and the conjecture:

Conjecture 4.9.2. *The satisfiability problem for $\mu^{<\infty}$-ML is decidable.*

is left open. Nevertheless, our automata model allows us to prove the conjecture for some special cases. The most advanced one is the satisfiability of the Büchi fragment over infinite words.

Apart from decidability, the relations between $\mu^{<\infty}$-ML and different logics deserve a deeper study. An interesting (although not included in this dissertation nor published) result identifies the fragment of μ-ML without nesting of the countdown operators with a certain multi-valued modal fixpoint logic. The natural comparison with other (un)boundedness-related logics: (fragments of) $\mathsf{MSO} + \mathsf{U}$, $\mathsf{WMSO} + \mathsf{U}$ and cost logics, is less straightforward than it could seem and requires further study.

Another question is whether our results generalize to structures of different *type* than a Kripke model: ordered trees, weighted graphs etc. Here the answer is simple and positive: under mild assumptions the results generalize. This can be formally expressed using category-theoretic notion of a coalgebra, which offers a uniform point of view on such structures of various types. The classical logic and automata have been successfully lifted to the coalgebraic setting (see e.g. [36]). Interestingly, this extension of type turns out to be orthogonal to our lifting of the classical constructions to the countdown setting. The combination of these two extensions into *coalgebraic countdown* logic and automata does not require any new ideas. In particular, nearly the entire Chapter 4 could be rewritten into the coalgebraic framework.

Finally, due to the abstract nature of the countdown operators they are well-defined in every complete lattice. Since complete lattices and fixpoints are ubiquitous in logic and computer science, this opens a wide range of new, possibly fruitful directions for research.

Bibliography

[1] Hans Bekić. Definable operation in general algebras, and the theory of automata and flowcharts. In *Programming Languages and Their Definition*, 1984.

[2] Patrick Blackburn, Maarten de Rijke, and Yde Venema. *Modal Logic*. Cambridge University Press, 2002.

[3] Mikołaj Bojańczyk. Weak mso with the unbounding quantifier. *Theory of Computing Systems*, 48(3):554–576, 2011.

[4] Mikołaj Bojańczyk and Tomasz Idziaszek. Algebra for infinite forests with an application to the temporal logic EF. In *CONCUR*, volume 5710 of *Lecture Notes in Computer Science*, pages 131–145. Springer, 2009.

[5] Mikołaj Bojańczyk, Edon Kelmendi, Rafał Stefański, and Georg Zetzsche. Extensions of ω-regular languages. In *Proceedings of the 35th Annual ACM/IEEE Symposium on Logic in Computer Science*, LICS '20, page 266–272, 2020.

[6] Mikolaj Bojanczyk, Pawel Parys, and Szymon Torunczyk. The MSO+U Theory of (N,<) Is Undecidable. In *33rd Symposium on Theoretical Aspects of Computer Science (STACS 2016)*, volume 47 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 21:1–21:8, 2016.

[7] Mikołaj Bojanczyk and Szymon Torunczyk. Weak MSO+U over infinite trees. In *29th International Symposium on Theoretical Aspects of Computer Science (STACS 2012)*, volume 14 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 648–660, 2012.

[8] Mikołaj Bojańczyk. A bounding quantifier. In Jerzy Marcinkowski and Andrzej Tarlecki, editors, *Computer Science Logic*, pages 41–55, 09 2004.

[9] Mikołaj Bojańczyk and Thomas Colcombet. Bounds in w-regularity. In *Proceedings - Symposium on Logic in Computer Science*, pages 285 – 296, 01 2006.

[10] Mikołaj Bojańczyk, Laure Daviaud, Bruno Guillon, Vincent Penelle, and A. V. Sreejith. Undecidability of mso+"ultimately periodic". *ArXiv*, abs/1807.08506, 2018.

[11] Cristian S. Calude, Sanjay Jain, Bakhadyr Khoussainov, Wei Li, and Frank Stephan. Deciding parity games in quasi-polynomial time. *SIAM Journal on Computing*, 0(0), 2017.

[12] Thomas Colcombet. Regular Cost Functions, Part I: Logic and Algebra over Words. *Logical Methods in Computer Science*, Volume 9, Issue 3, 2013.

[13] Anuj Dawar and Martin Otto. Modal characterisation theorems over special classes of frames. *20th Annual IEEE Symposium on Logic in Computer Science (LICS' 05)*, pages 21–30, 2005.

[14] Erwin Engeler. A characterization of theories with isomorphic denumerable models. *Notices of the American Mathematical Society*, 6, 1959.

[15] Lauri Hella, Antti Kuusisto, and Raine Rönnholm. Bounded game-theoretic semantics for modal mu-calculus. In *GandALF*, 2020.

[16] Matthew Hennessy and Robin Milner. Algebraic laws for non-determinism and concurrency. *Journal of the ACM*, 32:137–161, 1985.

[17] Graham Higman. Ordering by divisibility in abstract algebras. *Proceedings of The London Mathematical Society*, pages 326–336, 1952.

[18] Szczepan Hummel and Michał Skrzypczak. The topological complexity of mso+u and related automata models. *Fundamenta Informaticae*, 119:87–111, 2012.

[19] Bart Jacobs and Jan J. M. M. Rutten. *An introduction to (co)algebra and (co)induction*, pages 38–99. Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, 2011.

[20] David Janin and Igor Walukiewicz. On the expressive completeness of the propositional mu-calculus with respect to monadic second order logic. In Ugo Montanari and Vladimiro Sassone, editors, *CONCUR '96: Concurrency Theory*, pages 263–277, 1996.

[21] Jędrzej Kołodziejski. Bisimulational categoricity. Master's thesis, University of Warsaw, 2018.

[22] Jędrzej Kołodziejski. Bisimulational categoricity. In Nicola Olivetti, Rineke Verbrugge, Sara Negri, and Gabriel Sandu, editors, *13th Conference on Advances in Modal Logic, AiML 2020, Helsinki, Finland, August 24-28, 2020*, pages 389–404. College Publications, 2020.

[23] Jędrzej Kołodziejski and Bartek Klin. Countdown μ-Calculus. In Stefan Szeider, Robert Ganian, and Alexandra Silva, editors, *47th International Symposium on Mathematical Foundations of Computer Science (MFCS 2022)*, volume 241 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 64:1–64:14, Dagstuhl, Germany, 2022. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.

[24] David Marker. *Model Theory: An Introduction*. Graduate Texts in Mathematics. Springer, 2013.

[25] Wojciech Czerwiński Mikołaj Bojańczyk. *Automata Toolbox*. Univrsity of Warsaw, 2018.

[26] Jan J. M. M. Rutten. Universal coalgebra: a theory of systems. *Theor. Comput. Sci.*, 249:3–80, 2000.

[27] Czesław Ryll-Nardzewski. On the categoricity in power $\aleph_0$. *Bulletin de l'Académie Polonaise des Sciences, Série des Sciences Mathématiques, Astronomiques Et Physiques*, 7:545–548, 1959.

[28] Davide Sangiorgi. On the origins of bisimulation and coinduction. *ACM Trans. Program. Lang. Syst.*, 31(4):15:1–15:41, 2009.

[29] Davide Sangiorgi. *Introduction to Bisimulation and Coinduction*. Cambridge University Press, 2012.

[30] Michał Skrzypczak. *Descriptive Set Theoretic Methods in Automata Theory - Decidability and Topological Complexity*, volume 9802 of *Lecture Notes in Computer Science*. Springer, 2016.

[31] Lars Svenonius. $\aleph_0$-categoricity in first-order predicate calculus. *Theoria (Lund)*, 25:82–94, 1959.

[32] Alfred Tarski. A lattice-theoretical fixpoint theorem and its applications. *Pacific Journal of Mathematics*, 5:285–309, 1955.

[33] Johan van Benthem. *Modal Correspondence Theory*. PhD thesis, University of Amsterdam, 1977.

[34] Johan van Benthemm. Modal logic: A contemporary view. *Internet Encyclopedia of Philosophy* (`https://iep.utm.edu/modal-lo`, 31.07.2023).

[35] Rob van Glabbeek. The linear time-branching time spectrum. In J. C. M. Baeten and J. W. Klop, editors, *CONCUR '90 Theories of Concurrency: Unification and Extension*, pages 278–297, Berlin, Heidelberg, 1990. Springer Berlin Heidelberg.

[36] Yde Venema. Automata and fixed point logic: A coalgebraic perspective. *Inf. Comput.*, 204:637–678, 2006.

[37] Yde Venema. Lectures on the modal μ-calculus, 2020.

[38] Igor Walukiewicz. Pushdown processes: Games and model-checking. *Information and Computation*, 164(2):234–263, 2001.