

Generalized Equivalence Checking of Concurrent Programs

Benjamin Bisping

<https://bbisping.de>

Technische Universität Berlin

Table of contents

1	Introduction: What's the Difference?	1
1.1	Linear-Time–Branching-Time Spectroscopy	1
1.2	This Thesis	2
1.3	Artifacts and Papers	2
2	Preliminaries: Communicating Systems and Games	6
2.1	Behavior of Programs	7
2.1.1	Labeled Transition Systems	7
2.1.2	Calculus of Communicating Systems	9
2.2	Behavioral Equivalences	11
2.2.1	Trace Equivalence	11
2.2.2	Simulation and Bisimulation	12
2.2.3	The Hierarchy of Equivalences	13
2.2.4	Congruences	14
2.3	Modal Logic	15
2.3.1	Hennessy–Milner Logic to Express Observations	15
2.3.2	Characterizing Bisimilarity via HML	16
2.3.3	The Perks of Modal Characterizations	17
2.3.4	Expressiveness and Distinctiveness	18
2.4	Games	19
2.4.1	Reachability Games	19
2.4.2	The Semantic Game of HML	21
2.4.3	The Bisimulation Game	22
2.4.4	How Bisimulation Game and HML Are Connected	25
2.4.5	Deciding Reachability Games	27
2.5	Discussion	28
3	Context: The Spectrum of Equivalences	30
3.1	Observability Hierarchies	30
3.1.1	Understanding the Equivalence Hierarchy through Modal Logics	31
3.1.2	Incomparabilities	31
3.1.3	Lattices	32
3.2	The Linear-Time–Branching-Time Spectrum	33
3.2.1	Spectra as Observability Lattices	33
3.2.2	The Strong Notions of Observability	35
3.2.3	The Strong Linear-Time–Branching-Time Spectrum	37
3.2.4	Non-Intersectionality	39
3.3	Spectroscopy	40
3.3.1	The Spectroscopy Problem	40
3.3.2	Spectroscopy as Abstract Subtraction	42
3.3.3	Complexities	42

3.4	Discussion	43
4	Approach: Equivalence Problems as Energy Games	45
4.1	Energy Games	46
4.1.1	Monotonic Energy Games	46
4.1.2	Declining Energy Games	49
4.2	Characterizing the P-Time Part of the Spectrum	50
4.2.1	The Polynomial Slice	50
4.2.2	The Bisimulation <i>Energy Game</i>	52
4.2.3	Correctness of Characterization	53
4.3	Deciding Energy Games	56
4.3.1	Galois Connections	56
4.3.2	The Algorithm	57
4.3.3	Complexity	60
4.3.4	Solving Declining Energy Games	61
4.3.5	Polynomial Spectroscopy Complexity	63
4.4	Discussion	64
5	Spectroscopy of the Strong Equivalence Spectrum	66
5.1	The Strong Spectroscopy Game	66
5.1.1	The Game	66
5.1.2	Correctness	67
5.1.3	Complexity	67
5.2	More Specific Games	67
5.2.1	The Clever Game	67
5.2.2	Instance Games	67
5.3	Discussion	67
6	Spectroscopy for the Weak Silent-Step Spectrum	68
6.1	Lifting to the Spectrum of Stable Equivalences	68
6.1.1	Stable Behavioral Equivalences	68
6.1.2	The Adapted Spectroscopy Game	68
6.2	The Weak Spectrum	68
6.2.1	Special Weak Equivalences	68
6.2.2	HML with Silent Behavior	68
6.2.3	Expressing the Weak Spectrum with Quantities	68
6.3	The Weak Spectroscopy Game	68
6.3.1	The Game	68
6.3.2	Correctness	68
6.4	Discussion	68
7	Implementations	69
7.1	Prototype Implementation: <i>equiv.io</i>	69
7.2	GPU Implementation: <i>gpuequiv</i>	69
7.3	Other Implementations	69
7.3.1	Computer Game: <i>The Spectroscopy Invaders</i>	69
7.3.2	CAAL Extension	69
8	Conclusion	70

Table of contents

iii

References

71

Read this Book! (Once It's Finished)

Warning

This is a work-in-progress document. Handle with care!

You should read this book (once it's finished) if you're interested in the semantics of concurrent programs and want to understand ...

- How behavioral process equivalences are treated uniformly in equivalence spectra.
- How modal-logical and relational characterizations of behavioral equivalences are actually just “shadows” of game characterizations.
- How energy games can be used to turn a family of qualitative equivalence problems into one quantitative problems.
- How all behavioral equivalences can be decided at once.

TODO

(This is a space for TODOs that do not yet have a position in the text body.)

1 Introduction: What's the Difference?

Have you ever looked at two program models and wondered *how equivalent they are* —or, conversely: how they can be distinguished?

I have run into this problem often, for instance when analyzing models of distributed algorithms or when devising examples for teaching. But obviously, the question already occurs every time one rewrites a program part and hopes for it to still do its job.

The first time I had to formally face the question was when working as a student research assistant: I was tasked with implementing a translation from the process algebra Timed CSP to Timed Automata.¹ ... How to tell whether the translation would properly honor the semantics of the two formalisms? Did it translate CSP terms to automata *with the same meaning*? Even the definition of the question is tricky, as there are different notions of what counts as “same meaning” in the semantics of programs.

I then took my very first look into the seminal paper on the landscape of process equivalences with internal behavior, the “Linear-time–branching-time spectrum II” by van Glabbeek (1993). Its central figure, reproduced in Figure 1.1, mesmerized me. But, how to use the spectrum to decide for a given pair of processes which of the many equivalences apply, is far from straightforward. Over the years, I have learned that others too have run into this problem: For instance, Nestmann & Pierce (2000) thinking about process algebra encodings and Bell (2013) verifying compiler optimizations. The problem can be summarized as follows:

How does one conveniently decide for a pair of systems which notions from a spectrum of behavioral equivalences equate the two?

Above question will be the *research question of this thesis*. We² want to enable future researchers to tap into the wisdom of the linear-time–branching-time spectrum and to easily determine what equivalences fit their models.

1.1 Linear-Time–Branching-Time Spectroscopy

Van Glabbeek’s papers on comparative concurrency semantics (1990, 1993) treat a zoo of distinct *qualitative questions* of the form “Are processes p and q equivalent with respect to notion N ?”, where N would for example be trace




Figure 1.1: The linear-time–branching-time spectrum with silent moves as depicted in Glabbeek (1993).

¹ The aim was to bridge between the tools FDR2 and UPPAAL for Göthel’s dissertation (2012). Previous work I was to base the translations on had serious flaws: One approach introduced spurious deadlocks to the model, the other was unable to handle nesting of choices and parallel composition. Clearly, we had to change the encoding!

² As you might have noticed, this text uses more personal pronouns than is common in much of computer science literature. Their meaning is the same as in a lecture: “I” = “The author”, “You” = “The audience”, “We” = “The asynchronous collective of author and readers.”

or bisimulation equivalence. He unveils an underlying structure where the questions can easily be compared with respect to their distinctive power. This is analogous to the *spectrum of light* where seemingly qualitative properties (“The light is blue / green / red.”) happen to be *quantitative* (“The distribution of wavelengths peaks at 460/550/630 nm.”).

For light (i.e. electromagnetic radiation), the mix of wavelengths can be determined through a process called *spectroscopy*. So, we may ask:

If there exists a “linear-time–branching time spectrum,” does this mean, there also is some kind of “linear-time–branching-time spectroscopy”?

This thesis answers the question positively, which is the key step to tackle our research question. One can compute what mix of (in-)distinguishabilities exists between a pair of processes, and this places the two on the spectrum of equivalences. We thus turn a set of qualitative equivalence problems into one quantitative problem of *how equivalent* two systems are. As we will see, this amounts to an abstract form of subtraction between programs, determining what kinds of differences an outside examiner might observe.

1.2 This Thesis

The algorithm at the core of this thesis describes how to *decide all behavioral equivalences at once* for varying spectra of equivalence using *games that can count to limit possible distinctions by attacker energy levels*. More precisely, we make the following contributions:

- Chapter 2 lays some foundations and makes precise *how bisimilarity games relate to grammars of distinguishing formulas* of Hennessy–Milner modal logic.
- Chapter 3 shows *how to understand the strong linear-time–branching-time spectrum quantitatively* and formalizes the *spectroscopy problem*.
- Chapter 4 introduces the approach of *characterizing equivalence spectra through a class of energy games* and *how to decide such games*.
- TODO Chapter 5 applies the approach to decide the whole strong spectrum through one *game for linear-time–branching-time spectroscopy*.
- TODO Chapter 6 adapt the game for the *weak spectrum of equivalences accounting for silent-steps*.
- TODO Chapter 7 showcases implementations to conveniently perform equivalence spectroscopies in web browsers and using GPU parallelization.

1.3 Artifacts and Papers

This thesis ties together the work of several publications in a coherent presentation. It is written to be understandable on its own. For details, we usually

refer to the original publications or to other artifacts for implementation and machine-checked proofs.

Publications. The following four publications (with me as main author) fuel the following chapters:

- “Deciding All Behavioral Equivalences at Once: A Game for Linear-Time–Branching-Time Spectroscopy” (LMCS 2022) introduces the *spectroscopy problem* and the *core idea* to decide the whole strong spectrum using games that search trees of possible distinguishing formulas. (Conference version: TACAS 2021)
- “Process Equivalence Problems as Energy Games” (CAV 2023a) makes a *big technical leap by using energy games*, removing the necessity for explicit formula construction. (Tech report: arXiv 2023b)
- “Characterizing Contrasmilarity through Games, Modal Logic, and Complexity” (Information & Computation 2024) closes *gaps in the weak spectrum of equivalences* for games and their complexity and their link to modal logics. (Isabelle/HOL theory: AFP 2023; Workshop version: EXPRESS/SOS 2021)
- “One Energy Game for the Spectrum Between Branching Bisimilarity and Weak Trace Semantics” (EXPRESS/SOS 2024) adapts the *spectroscopy approach for the weak spectrum*.

Prototype. The algorithms of this thesis have been validated through a *Scala prototype implementation*. It can be tried out in the browser on <https://equiv.io>. The source is available openly on <https://github.com/benkeks/equivalence-fiddle>.

Isabelle formalization. Instead of swamping the document with technical proofs, these are contained in machine-checkable Isabelle/HOL theories.

- **Equivalence spectrum formalization** contains proofs for the early chapters: <https://benkeks.github.io/ltbt-spectroscopy-isabelle/>.
- **Weak Equivalence Spectroscopy** formalizes the core results of Chapter 6: <https://equivio.github.io/silent-step-spectroscopy>. The theory has been developed together with TU Berlin students Lisa Annett Barthel, Leonard Moritz Hübner, Caroline Lemke, Karl Parvis Philipp Mattes, and Lenard Mollenkopf.

Student theses. Some parts of this thesis strongly rely on student work that I have supervised, in particular, on the following theses:

- Trzeciakiewicz (2021): *Linear-Time–Branching-Time Spectroscopy as an Educational Web Browser Game* provides a computer game version of the spectroscopy procedure to be discussed in Subsection 7.3.1.
- Ozegowski (2023): *Integration eines generischen Äquivalenzprüfers in CAAL* extends the Concurrency Workbench Aalborg Edition with a spectroscopy feature, reported in Subsection 7.3.2.

- Mattes (2024): *Measuring Expressive Power of HML Formulas in Isabelle/HOL* proves the approach to modal characterization of Section 3.2.
- Vogel (2024): *Accelerating Process Equivalence Energy Games using WebGPU*, topic of Section 7.2, allows massive parallelization on the GPU of key parts of our algorithm.
- Lemke (2024): *A Formal Proof of Decidability of Multi-Weighted Declining Energy Games* formalizes the fixedpoint algorithm of Section 4.3.

Other publications. Only indirectly connected to this dissertation project are my prior publications in leading me to topic and techniques:

- Bisping et al. (2016, ITP): “Mechanical Verification of a Constructive Proof for FLP.”
- Bisping & Nestmann (2019, TACAS): “Computing Coupled Similarity.”
- Bisping et al. (2020, Acta Informatica): “Coupled similarity: the first 32 years.”

Other Bachelor theses. During my dissertation project, I also supervised several other Bachelor theses, many of which played important roles in shaping the research. Although they do not appear directly on the following pages, I want to acknowledge the students’ vital contribution to the research.

- Peacock (2020): *Process Equivalences as a Video Game*
- Lê (2020): *Implementing Coupled Similarity as an Automated Checker for mCRL2*
- Wrusch (2020): *Ein Computerspiel zum Erlernen von Verhaltensäquivalenzen*
- Reichert (2020): *Visualising and Model Checking Counterfactuals*
- Wittig (2020): *Charting the Jungle of Process Calculi Encodings*
- Bulik (2021): *Statically Analysing Inter-Process Communication in Elixir Programs*
- Montanari (2021): *Kontrasimulation als Spiel*
- Pohlmann (2021): *Reducing Strong Reactive Bisimilarity to Strong Bisimilarity*
- England (2021): *HML Synthesis of Distinguished Processes*
- Duong (2022): *Developing an Educational Tool for Linear-Time–Branching-Time Spectroscopy*
- Alshukairi (2022): *Automatisierte Reduktion von reaktiver zu starker Bisimilarität*
- Adler (2022): *Simulation fehlertoleranter Konsensalgorithmen in Hash.ai*
- Sandt (2022): *A Video Game about Reactive Bisimilarity*
- Lönne (2023): *An Educational Computer Game about Counterfactual Truth Conditions*
- Hauschild (2023): *Nonlinear Counterfactuals in Isabelle/HOL*
- Stöcker (2024): *Higher-Order Diadic μ -Calculus—An Efficient Framework for Checking Process Equivalences?*

- Kurzan (2024): *Implementierung eines Contrasimilarity-Checkers für mCRL2*

2 Preliminaries: Communicating Systems and Games

There is a virtually infinite supply of formalisms to model programs and communicating systems. They are accompanied by many notions for behavioral equivalence and refinement to relate or distinguish their models.

This chapter takes a tour into this field. The tour is *agnostic* in building on the most basic and versatile formalism of *labeled transition systems* and on standard equivalences such as *trace equivalence* and *bisimilarity*. The approach is also *opinionated* in focussing on Robin Milner’s tradition of concurrency theory with a strong pull towards *game characterizations*. The very core concepts are described in Section 2.1.

We will explore the trinity of relational, modal, and game-theoretic characterizations of behavioral equivalences. Figure 2.1 shows the scaffold of this section, instantiated to the classic notion of bisimilarity:

- Section 2.2 introduces (among other notions) bisimilarity through its *relational characterization* in terms of symmetric simulation relations.
- Section 2.3 treats the dual connection of bisimilarity and the distinguishing powers of a *modal logic*, known as the *Hennessy–Milner theorem*.
- Section 2.4 shows that both, the relations to validate bisimilarity and the modal formulas to falsify it, appear as certificates of strategies in a *reachability game* where an attacker tries to states apart and a defender tries to prevent this.

Readers who are familiar with the contents of Figure 2.1 can mostly skim through this chapter of preliminaries. The two core insights that this chapter intends to seed for coming chapters are:

i Idea 1: It’s all a game!

Equivalence games are a most versatile way to handle behavioral equivalences and obtain decision procedures. The Hennessy–Milner theorem appears as a shadow of the determinacy of the bisimulation reachability game.

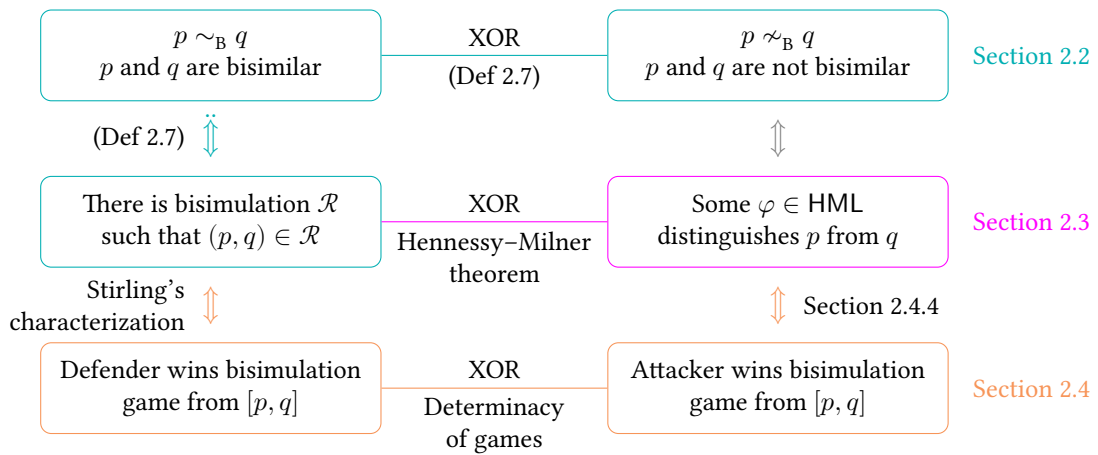


Figure 2.1: Core-correlations for bisimilarity between relational definition, modal distinguishability, and game.

i Idea 2: Modal first!

Modal characterizations allow a uniform handling of the hierarchy of equivalences. Productions in grammars of potential distinguishing formulas translate to game rules for equivalence games.

Both points might seem non-standard to those who are more used to relational or denotational definitions of behavioral equivalences. To provide them with a bridge, we will start with relational and denotational definitions—but once we have crossed the bridge, we will not go back.

2.1 Behavior of Programs

Every computer scientist has some model in their head of what it is that their algorithms and programs are doing. Usually these models are wrong,³ especially, once concurrency enters the picture. The area of *formal methods* tries to make models sufficiently precise that one at least can say what went wrong.

³ But some are useful, as the saying goes.

2.1.1 Labeled Transition Systems

Labeled transition systems are the standard formalism to discretely express the state space of programs, algorithms, and more. One can think of them as nondeterministic finite automata where the finiteness constraint has been lifted.

Definition 2.1 (Transition Systems). A *labeled transition system* (LTS) $\mathcal{S} = (\mathcal{P}, Act, \rightarrow)$ consists of

- \mathcal{P} , a set of *states*,

- *Act*, a set of *actions*, and
- $\rightarrow \subseteq \mathcal{P} \times \text{Act} \times \mathcal{P}$, a *transition relation*.⁴

We write $\text{Der}(p, \alpha)$ for the set of derivative states $\{p' \mid p \xrightarrow{\alpha} p'\}$, and $\text{Ini}(p)$ for the set of enabled actions $\{\alpha \mid \exists p'. p \xrightarrow{\alpha} p'\}$.

There is a canonical example used to discuss equivalences within transition systems, which we want to draw from. We will take the formulation that Henzinger used at CAV'23 as seen in Figure 2.2.

Example 2.1 (A Classic Example). Consider the transition system $\mathcal{S}_{PQ} = (\{P, p_a, p_b, p_1, p_2, Q, q_{ab}, q_1, q_2\}, \{a, b, \tau\}, \rightarrow_{PQ})$ given by the following graph:




Figure 2.3: Example system \mathcal{S}_{PQ} .

The program described by the transitions from P chooses non-deterministically during a τ -step between two options and then offers only either a or b . The program Q on the other hand performs a τ -step and then offers the choice between options a and b to the environment.

There are two things one might wonder about Example 2.1:

1. Should one care about non-determinism in programs? Subsection 2.1.2 shows how non-determinism arises naturally in concurrent programs.
2. Should one consider P and Q equivalent? This heavily depends. Section 2.2 will introduce a notion of equivalence under which the two are equivalent and one under which they differ.

Remark 2.1 (A Note on τ). The action τ (the greek letter “tau”) will in later chapters stand for *internal behavior* and receive special treatment. For the scope of this and the following three chapters, τ is an action like every other.

Generally, this thesis aims to be consistent with notation and naming in surrounding literature. Where there are different options, we usually prefer the less-greek one. Also, we will usually write literals and constant names in sans-serif and variables in *italics*. For the internal action, the whole field has converged to τ in *italics*, however—so, we will run with this.

⁴ locale `Labeled_Transition_Systems.lts`




Figure 2.2: Tom Henzinger employing Example 2.1 during CAV'23.

2.1.2 Calculus of Communicating Systems

To talk about programs in this thesis, we use Milner’s (1989) *Calculus of Communicating Systems* (CCS), which—together with other great contributions—earned him the Turing award. It is a tiny concurrent programming language that fits in your pocket, and can be properly described mathematically!

Definition 2.2 (Calculus of Communicating Systems). Let \mathcal{A} be a set of channel names, and \mathcal{X} a set of process names. Then, CCS processes, communicating via actions $\text{Act}_{\text{CCS}} := \mathcal{A} \cup \{\bar{\alpha} \mid \alpha \in \mathcal{A}\} \cup \{\tau\}$, are given by the following grammar:

$P ::=$	$\alpha.P$	with $\alpha \in \mathcal{A}$	“input action prefix”
	$\bar{\alpha}.P$	with $\alpha \in \mathcal{A}$	“output action prefix”
	$\tau.P$		“internal action”
	0		“null process”
	X	with $X \in \mathcal{X}$	“recursion”
	$P + P$		“choice”
	$P \mid P$		“parallel composition”
	$P \setminus A$	with $A \subseteq \mathcal{A}$	“restriction”

We call pairs of actions α and $\bar{\alpha}$ *coactions*, and work with the aliasing $\bar{\bar{\alpha}} = \alpha$. The intuition is that an action α represents receiving and $\bar{\alpha}$ expresses sending in communication. A pair of action and coaction can “react” in a communication situation and only become internal activity τ in the view of the environment.

Each sub-process tree must end in a 0 -process or recursion. For brevity, we usually drop a final 0 when writing terms, e.g., just writing ac for $ac.0$.

We place parenthesis, (...), in terms where the syntax trees are otherwise ambiguous, but understand the choice operator $+$ and the parallel operator \mid to be associative.


Example 2.2 (Concurrent Philosophers). Following tradition, we will express our examples in terms of philosophers who need forks to eat spaghetti.⁵ So, consider two philosophers P_A and P_B who want to grab a resource fork modeled as an action in order to eat where we express P_A eating with action a and P_B eating with b . The philosopher processes read:

$$\begin{aligned} P_A &:= \text{fork}.a.0 \\ P_B &:= \text{fork}.b.0 \end{aligned}$$

An LTS representation of P_A ’s behavior can be seen in the margin. Process P captures the whole scenario where the two philosophers compete for the fork using communication:

$$P := (\overline{\text{fork}}.0 \mid P_A \mid P_B) \setminus \{\text{fork}\}$$

⁵ Of course, you can just as well read the examples to be about computer programs that race for resources.



The restriction $\dots \setminus \{\text{fork}\}$ expresses that the fork-channel can only be used for communication within the system.

As the $\overline{\text{fork}}$ -action can be consumed by just one of the two philosophers, process P expresses exactly the program behavior seen in state P of Example 2.1.

The formal relationship between process terms and their LTS semantics is given by the following definition.

Definition 2.3 (CCS Semantics). Given an assignment of names to processes, $\mathcal{V}: \mathcal{X} \rightarrow \text{CCS}$, the operational semantics $\rightarrow_{\text{CCS}} \subseteq \text{CCS} \times \text{Act}_{\text{CCS}} \times \text{CCS}$ is defined inductively by the rules:

$$\begin{array}{c} \frac{}{\alpha.P \xrightarrow{\alpha}_{\text{CCS}} P} \quad \frac{P \xrightarrow{\alpha}_{\text{CCS}} P' \quad \mathcal{V}(X) = P}{X \xrightarrow{\alpha}_{\text{CCS}} P'} \\ \\ \frac{P_1 \xrightarrow{\alpha}_{\text{CCS}} P'_1}{P_1 + P_2 \xrightarrow{\alpha}_{\text{CCS}} P'_1} \quad \frac{P_2 \xrightarrow{\alpha}_{\text{CCS}} P'_2}{P_1 + P_2 \xrightarrow{\alpha}_{\text{CCS}} P'_2} \\ \\ \frac{P_1 \xrightarrow{\alpha}_{\text{CCS}} P'_1}{P_1 | P_2 \xrightarrow{\alpha}_{\text{CCS}} P'_1 | P_2} \quad \frac{P_2 \xrightarrow{\alpha}_{\text{CCS}} P'_2}{P_1 | P_2 \xrightarrow{\alpha}_{\text{CCS}} P_1 | P'_2} \\ \\ \frac{P_1 \xrightarrow{\alpha}_{\text{CCS}} P'_1 \quad P_2 \xrightarrow{\bar{\alpha}}_{\text{CCS}} P'_2}{P_1 | P_2 \xrightarrow{\tau}_{\text{CCS}} P'_1 | P'_2} \quad \frac{P \xrightarrow{\alpha}_{\text{CCS}} P' \quad \alpha, \bar{\alpha} \notin A}{P \setminus A \xrightarrow{\alpha}_{\text{CCS}} P' \setminus A} \end{array}$$

A process $P \in \text{CCS}$ now denotes a position in the transition system $(\text{CCS}, \text{Act}_{\text{CCS}}, \rightarrow_{\text{CCS}})$ defined through Definition 2.2.

Feel free to go ahead and check that the transitions of Example 2.1 indeed match those that Definition 2.3 prescribes for P of Example 2.2! (For readability, Example 2.1, has shorter state names, however.) For instance, the transition $P \xrightarrow{\tau} p_a$ of Example 2.1 would be justified as follows:

$$\frac{\frac{\frac{\overline{\text{fork.a}} \xrightarrow{\text{fork}}_{\text{CCS}} a \quad \mathcal{V}(P_A) = \text{fork.a}}{P_A \xrightarrow{\text{fork}}_{\text{CCS}} a}}{\overline{\text{fork}} \xrightarrow{\text{fork}}_{\text{CCS}} 0 \quad P_A | P_B \xrightarrow{\text{fork}}_{\text{CCS}} a | P_B}}{\overline{\text{fork}} | P_A | P_B \xrightarrow{\tau}_{\text{CCS}} 0 | a | P_B} \quad \tau \notin \{\text{fork}\}}{\frac{(\overline{\text{fork}} | P_A | P_B) \setminus \{\text{fork}\} \xrightarrow{\tau}_{\text{CCS}} (0 | a | P_B) \setminus \{\text{fork}\}}{P \xrightarrow{\tau}_{\text{CCS}} (0 | a | P_B) \setminus \{\text{fork}\}} \quad \mathcal{V}(P) = (\overline{\text{fork}} | P_A | P_B) \setminus \{\text{fork}\}}$$

Non-determinism like in P of Example 2.1 can be understood as a natural phenomenon in models with concurrency. The model leaves unspecified which of two processes will consume an internal resource and, to the outsider, it is transparent which one took the resource until they communicate. There are other ways how non-determinism plays a crucial role in models, for instance, as consequence of abstraction or parts that are left open in specifications.

The second process Q of Example 2.1 can be understood as a deterministic sibling of P .

Example 2.3 (Deterministic Philosophers). A process matching the transitions from Q in Example 2.1 would be the following, where the philosophers take the fork as a team and then let the environment choose who of them eats:

$$Q := (\overline{\text{fork}} \mid \text{fork}.(a + b)) \setminus \{\text{fork}\}.$$

But are P and Q equivalent?

2.2 Behavioral Equivalences

A behavioral equivalence formally defines when to consider two processes (or states, or programs) as *equivalent*. Clearly, there might be different ways of choosing such a notion of equivalence. Also, sometimes we are interested in a behavioral *preorder*, for instance, as a way of saying that a program does “less than” what some specification allows.

This section will quickly introduce the most common representatives of behavioral equivalences: Trace equivalence (and preorder), simulation equivalence (and preorder), and bisimilarity. We will then observe that the notions themselves can be compared in a hierarchy of equivalences.

2.2.1 Trace Equivalence

Every computer science curriculum features automata and their languages sometime at the beginning. Accordingly, comparing two programs in terms of the sequences of input/output events they might expose, referred to as their *traces*, is a natural starting point to talk about behavioral equivalences.

Definition 2.4 (Traces). The set of traces of a state $\text{Traces}(p)$ is recursively defined as

- $() \in \text{Traces}(p)$,
- $\alpha \cdot \vec{w} \in \text{Traces}(p)$ if there is p' with $p \xrightarrow{\alpha} p'$ and $\vec{w} \in \text{Traces}(p')$.

Definition 2.5 (Trace Equivalence). Two states p and q are considered *trace-equivalent*, written $p \sim_T q$, if $\text{Traces}(p) = \text{Traces}(q)$.

States are *trace-preordered*, $p \preceq_T q$, if $\text{Traces}(p) \subseteq \text{Traces}(q)$.

Example 2.4. The traces for the processes of Example 2.1 would be $\text{Traces}(P) = \{(), \tau, \tau a, \tau b\} = \text{Traces}(Q)$. Consequently, P and Q are trace-equivalent, $P \sim_T Q$.

As $\text{Traces}(p_a) = \{(), a\} \subseteq \{(), a, b\} = \text{Traces}(q_{ab})$, p_a is trace-preordered to q_{ab} , $p_a \preceq_T q_{ab}$. This ordering is strict, that is, $q_{ab} \not\preceq_T p_a$, due to $b \in \text{Traces}(q_{ab})$ but $b \notin \text{Traces}(p_a)$. We could say that trace b constitutes a *difference* between q_{ab} and p_a .

⁶ We denote the empty word by $()$.

⁷ abbreviation [Labeled_Transition_Systems.lts.traces](#)

⁸ abbreviation [Strong_Equivalences.lts.trace_equivalent](#)

⁹ definition [Strong_Equivalences.lts.trace_preordered](#)

Proposition 2.1. *The trace preorder \preceq_T is indeed a preorder (i.e., transitive and reflexive)¹⁰ and trace equivalence \sim_T is indeed an equivalence relation (i.e., transitive, reflexive, and moreover symmetric).¹¹*

¹⁰ lemma [Strong_Equivalences.lts.trace_preorder_transitive](#)

¹¹ lemma [Strong_Equivalences.lts.trace_equivalence_equiv](#)

Trace equivalence and preorder assign programs straightforward *denotational semantics*: The sets of traces they might expose. For many languages, these mathematical objects can be constructed directly from the expressions of the language. With the idea that the program text “denotes” its possible executions, the set of traces is called a “denotation” in this context. CCS, as we use it, follows another approach to semantics, namely the one of “operational semantics,” where the meaning of a program is in how it might execute.

There are several reasons why computer scientist did content semselves with trace equivalence when studying interactive systems. The core argument is that, in this context, one usually does not want to consider processes like P and Q to be equivalent: The two might interact differently with an environment. For instance, assume there is a user that wants to communicate $(\dots \mid a.success.0) \setminus \{a\}$. In interactions with Q , they will always reach success; with P , there is a possibility of ending up deadlocked in parallel with p_b , never achieving success.

2.2.2 Simulation and Bisimulation

The other big approach to behavioral equivalence of programs is the one of relating parts of their state spaces to one-another. The idea here is to identify which states of one program can be used to *simulate* the behavior of the another.

Definition 2.6 (Simulation). A relation on states, $\mathcal{R} \subseteq \mathcal{P} \times \mathcal{P}$, is called a *simulation* if, for each $(p, q) \in \mathcal{R}$ and $\alpha \in Act$ with $p \xrightarrow{\alpha} p'$ there is a q' with $q \xrightarrow{\alpha} q'$ and $(p', q') \in \mathcal{R}$.¹²

¹² definition [Strong_Equivalences.lts.simulation](#)

Definition 2.7 ((Bi-)similarity). Simulation preorder, simulation equivalence and bisimilarity are defined as follows:

- p is *simulated by* q , $p \preceq_S q$, iff there is a simulation \mathcal{R} with $(p, q) \in \mathcal{R}$.¹³
- p is *similar to* q , $p \sim_S q$, iff $p \preceq_S q$ and $q \preceq_S p$.¹⁴
- p is *bisimilar to* q , $p \sim_B q$, iff there is a *symmetric* simulation \mathcal{R} (i.e. $\mathcal{R} = \mathcal{R}^{-1}$) with $(p, q) \in \mathcal{R}$.¹⁵

¹³ definition [Strong_Equivalences.lts.simulated_by](#)

¹⁴ abbreviation [Strong_Equivalences.lts.similar](#)

¹⁵ definition [Strong_Equivalences.lts.bisimilar](#)

We also call a symmetric simulation *bisimulation* for short.¹⁶ Canceled symbols of relations refer to their negations, for instance, $p \not\sim_S q$ iff there is *no* simulation \mathcal{R} with $(p, q) \in \mathcal{R}$.

¹⁶ Other authors use a weaker definition, namely, that \mathcal{R} is a bisimulation if \mathcal{R} and \mathcal{R}^{-1} are simulations. Both definitions lead to the characterization of the same notion of bisimilarity.

Example 2.5. The following relations are simulations on the LTS of Example 2.1:

- the empty relation $\mathcal{R}_\emptyset := \emptyset$;

- the identity relation $\mathcal{R}_{\text{id}} := \text{id}_{\{P, P_a, P_b, P_1, P_2, Q, Q_{ab}, q_1, q_2\}} = \{(P, P), (P_a, P_a), (P_b, P_b), (P_1, P_1), (P_2, P_2), (Q, Q), (Q_{ab}, Q_{ab}), (q_1, q_1), (q_2, q_2)\}$;
- the universal relation between all final states $\mathcal{R}_{\text{fin}} := \{p_1, p_2, q_1, q_2\}^2$,
- more generally, the relation from final states to all other states: $\mathcal{R}_{\text{up}} := \{P_1, P_2, q_1, q_2\} \times \mathcal{P}$;
- a minimal simulation for P and Q: $\mathcal{R}_{\text{PQ}} := \{(P, Q), (P_a, q_{ab}), (P_b, q_{ab}), (P_1, q_1), (P_2, q_2)\}$;
- and the combination of the above $\mathcal{R}_{\text{max}} := \mathcal{R}_{\text{sim}} \cup \mathcal{R}_{\text{id}} \cup \mathcal{R}_{\text{up}}$.

The simulation \mathcal{R}_{PQ} shows that $P \preceq_S Q$.

However, there is no simulation that preorders Q to P, as there is no way to simulate the transition $Q \xrightarrow{a} q_{ab}$ from P for lack of a successor that allows a and b as does q_{ab} . (In Section 2.3, we will discuss how to capture such differences more formally.)

Thus, $Q \not\preceq_S P$, and $P \not\sim_S Q$. Moreover, there cannot be a symmetric simulation, $P \not\sim_B Q$.

Proposition 2.2. *The simulation preorder \preceq_S is indeed a preorder¹⁷, and \sim_S and \sim_B are equivalences.¹⁸*

2.2.3 The Hierarchy of Equivalences

Bisimilarity, similarity and trace equivalence form a small hierarchy of equivalences in the sense that they imply one-another in one direction, but not in the other. Let us quickly make this formal:

Lemma 2.1. *The relation \sim_B is itself a symmetric simulation.¹⁹*

Corollary 2.1. *If $p \sim_B q$, then $p \sim_S q$.²⁰*

Lemma 2.2. *If $p \preceq_S q$, then $p \preceq_T q$.²¹ (Consequently, $p \sim_S q$ also implies $p \sim_T q$.²²)*

We also have seen with example Example 2.5 that this hierarchy is strict between trace and simulation preorder in the sense that there exist p, q with $p \preceq_T q$ but not $p \preceq_S q$. The hierarchy also is strict between similarity and bisimilarity as the following example shows.

Example 2.6 (Trolled philosophers). Let us extend Q of Example 2.3 to include a troll process that might consume the fork and then do nothing:

$$T := (\overline{\text{fork}} \mid \text{fork} \mid \text{fork} \cdot (a + b)) \setminus \{\text{fork}\}.$$

This adds another deadlock state to the transition system, seen in Figure 2.5.

To similarity, this change is invisible, that is $Q \sim_S T$. (Reason: The relation $\{(Q, T), (T, Q)\} \cup \text{id}_{\{q_{ab}, q_1, q_2, q_3\}}$ is a simulation.)

However, to bisimilarity, $T \xrightarrow{\tau} q_3$ constitutes a difference. There cannot be a symmetric simulation handling this transition as Q has no deadlocked successors. Thus $Q \not\sim_B T$.

¹⁷ lemma [Strong_Equivalences.lts.simulation_preorder_transitive](#)

¹⁸ lemma [Strong_Equivalences.lts.bisimilarity_equiv](#)




Figure 2.4: Core hierarchy of equivalences.

¹⁹ lemma [Strong_Equivalences.lts.bisim_bisim](#)

²⁰ lemma [Strong_Equivalences.lts.bisim_bisim](#)

²¹ lemma [Strong_Equivalences.lts.sim_implies_trace_preord](#)

²² lemma [Strong_Equivalences.lts.sim_eq_implies_trace_eq](#)




Figure 2.5: Example with new deadlock q_3 .

The equivalences we have been discussed so far could also be understood as abstractions of an even finer equivalence, namely graph isomorphism:

Definition 2.8 (Graph Isomorphism). A bijective function $f: \mathcal{P} \rightarrow \mathcal{P}$ is called a *graph isomorphism* on a transition system if, for all p, p', α , the transition $p \xrightarrow{\alpha} p'$ exists if and only if the transition $f(p) \xrightarrow{\alpha} f(p')$ exists.²³

Two states p and q are considered *graph-isomorphism-equivalent*, $p \sim_{\text{ISO}} q$, iff there is a graph isomorphism f with $f(p) = q$.²⁴

Example 2.7. Consider the transition system in Figure 2.6. $p_{\text{even}} \sim_{\text{ISO}} p_{\text{odd}}$ because $f := \{p_{\text{even}} \mapsto p_{\text{odd}}, p_{\text{odd}} \mapsto p_{\text{even}}\}$ is a graph isomorphism.

Lemma 2.3. The relation \sim_{ISO} is itself a symmetric simulation and thus $p \sim_{\text{ISO}} q$ implies $p \sim_{\text{B}} q$.²⁵

Once again, the hierarchy is strict because of bisimilarity being less restricted in the matching of equivalent states.

Example 2.8. Consider the processes $P_1 := (\overline{\text{fork}} \mid \text{fork}) \setminus \{\text{fork}\}$ and $P_2 := (\overline{\text{fork}} \mid \text{fork} \mid \text{fork}) \setminus \{\text{fork}\}$. P_1 can transition to $(0 \mid 0) \setminus \{\text{fork}\}$, while P_2 has two options, namely $(0 \mid 0 \mid \text{fork}) \setminus \{\text{fork}\}$ and $(0 \mid \text{fork} \mid 0) \setminus \{\text{fork}\}$. All three reachable processes are deadlocks and thus isomorphic. But $P_1 \not\sim_{\text{ISO}} P_2$ because no bijection can connect the one successor of P_1 and the two of P_2 . However, $P_1 \sim_{\text{B}} P_2$, as bisimilarity is more relaxed.

2.2.4 Congruences

One of the prime quality criteria for behavioral equivalences is whether they form *congruences* with respect to fitting semantics or other important transformations. A congruence relates mathematical objects that can stand in for one-another in certain contexts, which, for instance, allows term rewriting. The concept is closely linked to another core notion of mathematics: monotonicity.

Definition 2.9 (Monotonicity and Congruence). An n -ary function $f: B_1 \times \dots \times B_n \rightarrow C$ is called *monotonic* with respect to a family of partial orders (B_k, \leq_k) and (C, \leq) iff, for all $b \in B_1 \times \dots \times B_n$ and $b' \in B_1 \times \dots \times B_n$, it is the case that $b_k \leq b'_k$ for all $k \leq n$ implies that $f(b) \leq f(b')$. We will usually encounter settings where all components use the same order $(B_1, \leq_1) = \dots = (B_n, \leq_n) = (C, \leq)$

The relation \leq is then referred to as a *precongruence* for f . If \leq moreover is symmetric (and thus an equivalence relation), then \leq is called a *congruence* for f .

Example 2.9 (Parity as Congruence). As a standard example for a congruence, consider the equivalence relation of equally odd numbers

$$\sim_{\text{odd}} := \{(m, n) \in \mathbb{N} \times \mathbb{N} \mid m \bmod 2 = n \bmod 2\}.$$

²³ definition [Strong_Equivalences.Its.isomorphism](#)

²⁴ definition [Strong_Equivalences.Its.is_isomorphic_to](#)




Figure 2.6: Transition system with an isomorphic cycle.

²⁵ lemma [Strong_Equivalences.Its.iso_implies_bisim](#)

For instance, $1 \sim_{\text{odd}} 3$ and $0 \sim_{\text{odd}} 42$, but not $42 \sim_{\text{odd}} 23$.

\sim_{odd} is a congruence for addition and multiplication. For instance, the sum of two odd numbers will always be even; the product of two odd numbers, will always be odd.

But \sim_{odd} is no congruence for integer division. For instance, $2 \sim_{\text{odd}} 4$, but $2/2 = 1 \not\sim_{\text{odd}} 2 = 4/2$.

Lemma 2.4. *The operators of CCS (Definition 2.2) are congruences for trace equivalence and bisimilarity on the CCS transition system (Definition 2.3).*

Proof. This is a well-established fact. TODO: Citation!

As an example, let us prove the congruence property for choice, $+$, on bisimilarity, \sim_{B} . For any CCS context, we assume $P \sim_{\text{B}} R$ and $Q \sim_{\text{B}} T$. We have to show that $P + Q \sim_{\text{B}} R + T$, that is due to symmetry, that for any α with $P + Q \xrightarrow{\alpha}_{\text{CCS}} PQ'$ there is RT' such that $R + T \xrightarrow{\alpha}_{\text{CCS}} RT'$ and $PQ' \sim_{\text{B}} RT'$. By the semantics of CCS, the first step must be either due to $P \xrightarrow{\alpha}_{\text{CCS}} PQ'$ or, analogously, $Q \xrightarrow{\alpha}_{\text{CCS}} PQ'$. Without loss of generality, we only consider $P \xrightarrow{\alpha}_{\text{CCS}} PQ'$. Due to $P \sim_{\text{B}} R$, this step implies some R' with $R \xrightarrow{\alpha}_{\text{CCS}} R'$ and $PQ' \sim_{\text{B}} R'$. As $R + T \xrightarrow{\alpha}_{\text{CCS}} R'$, we can take this R' to prove our goal. \square

2.3 Modal Logic

Modal logics are logics in which one can formalize how facts in one possible world relate to other possible worlds. In the computer science interpretation, the possible worlds are *program states*, and typical statements have the form: “If X happens during the execution, then Y will happen in the next step.”

We care about modal logics as they too can be used to characterize behavioral equivalences. In this section, we will show how this characterization works and argue that, for our purposes of comparative semantics, modal characterizations are a superb medium.

2.3.1 Hennessy–Milner Logic to Express Observations

Hennessy & Milner (1980) introduce the modal logic that is now commonly called *Hennessy–Milner logic* as a “little language for talking about programs.” The idea is that HML formulas express “experiments” or “tests” that an observer performs interacting with a program.

Definition 2.10 (Hennessy–Milner Logic). Formulas of *Hennessy–Milner logic* HML are given by the grammar:²⁶

$$\begin{array}{l} \varphi ::= \langle \alpha \rangle \varphi \quad \text{with } \alpha \in \text{Act} \quad \text{“observation”} \\ \quad | \bigwedge_{i \in I} \varphi_i \quad \text{with index set } I \quad \text{“conjunction”} \\ \quad | \neg \varphi \quad \text{“negation”} \end{array}$$

²⁶ datatype [Hennessy_Milner_Logic.hml_formula](#)

We also write conjunctions as sets $\bigwedge\{\varphi_1, \varphi_2, \dots\}$. The empty conjunction $\bigwedge \emptyset$ is denoted by \top and serves as the nil-element of HML syntax trees. We also usually omit them when writing down formulas, e.g., shortening $\langle a \rangle \langle b \rangle \top$ to $\langle a \rangle \langle b \rangle$.

The intuition behind HML is that it describes *what sequences of behavior* one may or may not *observe* of a system. Observations $\langle \alpha \rangle \dots$ are used to build up possible action sequences; conjunctions $\bigwedge\{\dots\}$ capture branching points in decision trees; and negations $\neg \dots$ describe impossibilities.

Definition 2.11 (HML semantics). The semantics of HML $\llbracket \cdot \rrbracket : \text{HML} \rightarrow 2^{\mathcal{P}}$ is defined recursively by:²⁷

$$\begin{aligned} \llbracket \langle \alpha \rangle \varphi \rrbracket &:= \{p \mid \exists p' \in \llbracket \varphi \rrbracket. p \xrightarrow{\alpha} p'\} \\ \llbracket \bigwedge_{i \in I} \varphi_i \rrbracket &:= \bigcap_{i \in I} \llbracket \varphi_i \rrbracket \\ \llbracket \neg \varphi \rrbracket &:= \mathcal{P} \setminus \llbracket \varphi \rrbracket \end{aligned}$$

²⁷ primrec [Hennessy_Milner_Logic.lts.satisfies](#)

Example 2.10. Let us consider some observations on the system of Example 2.1.

- $\llbracket \langle \tau \rangle \langle a \rangle \rrbracket = \{P, Q\}$ as both, P and Q, expose the trace τa ,
- $Q \in \llbracket \langle \tau \rangle \bigwedge\{\langle a \rangle, \langle b \rangle\} \rrbracket$, but $P \notin \llbracket \langle \tau \rangle \bigwedge\{\langle a \rangle, \langle b \rangle\} \rrbracket$.
- $P \in \llbracket \langle \tau \rangle \neg \langle a \rangle \rrbracket$, but $Q \notin \llbracket \langle \tau \rangle \neg \langle a \rangle \rrbracket$.

2.3.2 Characterizing Bisimilarity via HML

We can now add the middle layer of our overview graphic in Figure 2.1: That two states are bisimilar precisely if they cannot be told apart using HML formulas.

Definition 2.12 (Distinctions and Equivalences). We say that formula $\varphi \in \text{HML}$ *distinguishes* state p from state q if $p \in \llbracket \varphi \rrbracket$ but $q \notin \llbracket \varphi \rrbracket$.²⁸

²⁸ abbreviation [LTS_Semantics.lts.distinguishes](#)

We say a sublogic $\mathcal{O} \subseteq \text{HML}$ *preorders* state p to q , $p \preceq_{\mathcal{O}} q$, if no $\varphi \in \mathcal{O}$ is distinguishing p from q .²⁹ If the preordering goes in both directions, we say that p and q are equivalent with respect to sublogic \mathcal{O} , written $p \sim_{\mathcal{O}} q$.³⁰

²⁹ definition [LTS_Semantics.lts.preordered](#)

³⁰ definition [LTS_Semantics.lts.equivalent](#)

By this account, $\langle \tau \rangle \neg \langle a \rangle$ of Example 2.10 distinguishes P from Q. On the other hand, $\langle \tau \rangle \bigwedge\{\langle a \rangle, \langle b \rangle\}$ distinguishes Q from P. (The direction matters!) For instance, the sublogic $\{\langle \tau \rangle \langle a \rangle, \langle \tau \rangle \langle b \rangle\}$ preorders P and Q in both directions; so the two states are equivalent with respect to this logic.

Proposition 2.3. Consider an arbitrary HML sublogic $\mathcal{O} \subseteq \text{HML}$. Then, $\preceq_{\mathcal{O}}$ is a preorder, and $\sim_{\mathcal{O}}$ an equivalence relation.³¹

³¹ lemma [LTS_Semantics.lts.equivalent_equiv](#)

Lemma 2.5. Hennessy–Milner logic equivalence \sim_{HML} is a simulation relation.³²

³² lemma [HML_Spectrum.lts.hml_equiv_sim](#)

Proof. Assume it were not. Then there would need to be $p \sim_{\text{HML}} q$ with step $p \xrightarrow{\alpha} p'$, and no q' such that $q \xrightarrow{\alpha} q'$ and $p' \sim_{\text{HML}} q'$. So there would need to be a distinguishing formula $\varphi_{q'}$ for each q' that q can reach by an α -step.

Consider the formula $\varphi_\alpha := \langle \alpha \rangle \bigwedge_{q' \in \text{Der}(q, \alpha)} \varphi_{q'}$. It must be true for p and false for q , contradicting $p \sim_{\text{HML}} q$. \square

Lemma 2.6 (HML Bisimulation Invariance). *If $p \in \llbracket \varphi \rrbracket$ and $p \sim_{\text{B}} q$ then $q \in \llbracket \varphi \rrbracket$.*³³

³³ lemma [HML_Spectrum.lts.hml_bisim_invariant](#)

Proof. Induct over the structure of φ with arbitrary p and q .

- Case $p \in \llbracket \langle \alpha \rangle \varphi \rrbracket$. Thus there is $p' \in \llbracket \varphi \rrbracket$ with $p \xrightarrow{\alpha} p'$. Because \sim_{B} is a simulation according to Lemma 2.1, this implies q' with $p' \sim_{\text{B}} q'$. The induction hypothesis makes $p' \in \llbracket \varphi \rrbracket$ entail $q' \in \llbracket \varphi \rrbracket$ and thus $q \in \llbracket \langle \alpha \rangle \varphi \rrbracket$.
- Case $p \in \llbracket \bigwedge_{i \in I} \varphi_i \rrbracket$. The induction hypothesis on the φ_i directly leads to $q \in \llbracket \bigwedge_{i \in I} \varphi_i \rrbracket$.
- Case $p \in \llbracket \neg \varphi \rrbracket$. Symmetry of \sim_{B} according to Proposition 2.2, implies $q \sim_{\text{B}} p$. By induction hypothesis, $q \in \llbracket \varphi \rrbracket$ implies $p \in \llbracket \varphi \rrbracket$. So, using contraposition, the case implies $q \in \llbracket \neg \varphi \rrbracket$. \square

Combining the bisimulation invariance for one direction and that \sim_{HML} is a symmetric simulation (Proposition 2.3 and Lemma 2.5) for the other, we obtain that HML precisely characterizes bisimilarity:

Theorem 2.1 (Hennessy–Milner Theorem). *Bisimilarity and HML equivalence coincide, that is, $p \sim_{\text{B}} q$ precisely if $p \sim_{\text{HML}} q$.*³⁴

³⁴ theorem [HML_Spectrum.lts.Hennessy_Milner_theorem](#)

Remark 2.2. In the standard presentation of the theorem, image finiteness of the transition system is assumed. This means that $\text{Der}(p, \alpha)$ is finite for every $p \in \mathcal{P}$. The amount of outgoing transitions matters precisely in the construction of φ_α in the proof of Lemma 2.5. But as our definition of HML (Definition 2.10) allows infinite conjunctions $\bigwedge_{i \in I} \dots$, we do not need an assumption here. The implicit assumption is that the cardinality of index sets I can match that of $\text{Der}(p, \alpha)$. The original proof by Hennessy & Milner (1980) uses binary conjunction ($\varphi_1 \wedge \varphi_2$) and thus can only express finitary conjunctions.

2.3.3 The Perks of Modal Characterizations

There is merit in also characterizing other equivalences through sublogics $\mathcal{O} \subseteq \text{HML}$. There are three immediate big advantages to modal characterization:

Modal characterizations lead to *proper preorders and equivalences by design* due to Proposition 2.3. That is, if a behavioral preorder (or equivalence) is defined through modal logics, there is no need of proving transitivity and reflexivity (and symmetry).

Secondly, modal characterizations can directly *unveil the hierarchy between preorders*, if defined cleverly, because of the following property.

Proposition 2.4. *If $\mathcal{O} \subseteq \mathcal{O}'$ then $p \preceq_{\mathcal{O}'} q$ implies $p \preceq_{\mathcal{O}} q$ for all p, q .*³⁵

³⁵ lemma [LTS_Semantics.lts_semantics_preorder_contraposition](#)

Pay attention to the opposing directions of \subseteq and implication, here!

Thirdly, as a corollary of Proposition 2.4, modal characterizations *ensure equivalences to be abstractions of bisimilarity*, which is a sensible base notion of equivalence.³⁶

In Chapter 3, we will discuss how the hierarchy of behavioral equivalences can be understood nicely and uniformly if viewed through the modal lense.

³⁶ Among other things, bisimilarity checking has a better time complexity than other notions as will be discussed in Subsection 3.3.3.

2.3.4 Expressiveness and Distinctiveness

Proposition 2.4 actually is a weak version of another proposition about *distinctiveness* of logics.

Definition 2.13. We say that an *Act*-observation language \mathcal{O} is *less or equal in expressiveness* to another \mathcal{O}' iff, for any *Act*-transition system, for each $\varphi \in \mathcal{O}$, there is $\varphi' \in \mathcal{O}'$ such that $\llbracket \varphi \rrbracket = \llbracket \varphi' \rrbracket$. (The definition can also be read with regards to a fixed transition system \mathcal{S} .)³⁷ If the inclusion holds in both directions, \mathcal{O} and \mathcal{O}' are *equally expressive*.

³⁷ definition [LTS_Semantics.lts_semantics.leq_expressive](#)

Definition 2.14. We say that an *Act*-observation language \mathcal{O} is *less or equal in distinctiveness* to another \mathcal{O}' iff, for any *Act*-transition system and states p and q , for each $\varphi \in \mathcal{O}$ that distinguishes p from q , there is $\varphi' \in \mathcal{O}'$ that distinguishes p from q .³⁸ If the inclusion holds in both directions, \mathcal{O} and \mathcal{O}' are *equally distinctive*.

³⁸ definition [LTS_Semantics.lts_semantics.leq_distinctive](#)

Lemma 2.7. *Subset relationship entails expressiveness entails distinctiveness.*

- If $\mathcal{O} \subseteq \mathcal{O}'$, then \mathcal{O} is less or equal in expressiveness to \mathcal{O}' .³⁹
- If \mathcal{O} is less or equal in expressiveness to \mathcal{O}' , then \mathcal{O} is less or equal in distinctiveness to \mathcal{O}' .⁴⁰

³⁹ lemma [LTS_Semantics.lts_semantics.subset_expressiveness](#)

⁴⁰ lemma [LTS_Semantics.lts_semantics.expressiveness_entails_distinctiveness](#)

The other direction does not hold. For instance, $\text{HML} \not\subseteq \text{HML} \setminus \{\top\}$, but they are equally expressive as $\neg\neg\top$ can cover for the removed element. At the same time, $\{\top\}$ is more expressive than \emptyset , but equally distinctive.

The stronger version of Proposition 2.4 is thus:

Proposition 2.5. *If \mathcal{O} is less or equal in distinctiveness to \mathcal{O}' then $p \preceq_{\mathcal{O}'} q$ implies $p \preceq_{\mathcal{O}} q$ for all p, q .*⁴¹

⁴¹ lemma [LTS_Semantics.lts_semantics.preorder_expressiveness_contraposition](#)

Often, an equivalence may be characterized by different sublogics. In particular, one may find smaller characterizations as in the following example for bisimilarity.

Example 2.11. Consider $\mathcal{O}_{[B]} \subseteq \text{HML}$ described by the following grammar.

$$\varphi^{[B]} ::= \langle \alpha \rangle \bigwedge_{i \in I} \varphi_i^{[B]} \quad | \quad \neg \varphi^{[B]}$$

$\mathcal{O}_{[B]}$ is a proper subset of HML. For instance, it lacks the observation $\langle a \rangle \langle b \rangle \top$. Due to the subset relation, $\mathcal{O}_{[B]}$ must be less or equal in expressiveness to HML, but this inclusion too is strict as \top cannot be covered for. But both logics are equally distinctive!

Lemma 2.8. $\mathcal{O}_{[\text{B}]}$ and HML are equally distinctive.⁴²

⁴² theorem [HML_Spectrum.lts.hml_and_bisimulation_game_observations_equally_expressive](#)

Proof. One direction is immediate from Lemma 2.7 as $\mathcal{O}_{[\text{B}]} \subseteq \text{HML}$.

For the other direction, we need to establish that for each $\varphi \in \text{HML}$ distinguishing some p from some q , there is a $\varphi' \in \mathcal{O}_{[\text{B}]}$ distinguishing p from q . We induct on the structure of φ with arbitrary p and q .

- Case $\langle \alpha \rangle \varphi$ distinguishes p from q . Thus there is p' with $p \xrightarrow{\alpha} p'$ distinguishing p' from all $q' \in \text{Der}(q, \alpha)$. By induction hypothesis, there must be $\varphi'_{q'} \in \mathcal{O}_{[\text{B}]}$ distinguishing p' from q' for each q' . Thus $\langle \alpha \rangle \bigwedge_{q' \in \text{Der}(q, \alpha)} \varphi'_{q'} \in \mathcal{O}_{[\text{B}]}$ distinguishes p from q .
- Case $\bigwedge_{i \in I} \varphi_i$ distinguishes p from q . Therefore some φ_i already distinguishes p from q . By induction hypothesis, there must be $\varphi'_i \in \mathcal{O}_{[\text{B}]}$ distinguishing p from q .
- Case $\neg \varphi$ distinguishes p from q . Thus φ distinguishes q from p . By induction hypothesis, there is $\varphi' \in \mathcal{O}_{[\text{B}]}$ distinguishing q from p . Accordingly, $\neg \varphi' \in \mathcal{O}_{[\text{B}]}$ distinguishes p from q . \square

The smaller bisimulation logic $\mathcal{O}_{[\text{B}]}$ will become important again later in the next section (in Subsection 2.4.4).

2.4 Games

So far, we have only seen behavioral equivalences and modal formulas as mathematical objects and not cared about decision procedures. This section introduces *game-theoretic characterizations* as a way of easily providing decision procedures for equivalences and logics alike. Intuitively, the games can be understood as dialogs between a party that tries to defend a claim and a party that tries to attack it.

2.4.1 Reachability Games

We use *Gale–Stewart-style reachability games* (in the tradition of [Gale & Stewart, 1953](#)) where the defender wins all infinite plays.

Definition 2.15 (Reachability Game). A *reachability game* $\mathcal{G} = (G, G_d, \mapsto)$ is played on a directed graph consisting of

- a set of *game positions* G , partitioned into
 - *defender positions* $G_d \subseteq G$
 - and *attacker positions* $G_a := G \setminus G_d$,
- and a set of *game moves* $\mapsto \subseteq G \times G$.⁴³

⁴³ locale [Equivalence_Games.game](#)

We denote by $\mathcal{G}(g_0)$ the game played from starting position $g_0 \in G$.

Definition 2.16 (Plays and Wins). We call the paths $g_0g_1\dots \in G^\infty$ with $g_i \rightsquigarrow g_{i+1}$ plays of $\mathcal{G}(g_0)$. They may be finite or infinite. The defender *wins* infinite plays. If a finite play $g_0\dots g_n \rightsquigarrow$ is stuck, the stuck player loses: The defender wins if $g_n \in G_a$, and the attacker wins if $g_n \in G_d$.

Usually, games are non-deterministic, that is, players have choices how a play proceeds at their positions. The player choices are formalized by *strategies*:

Definition 2.17 (Strategies and Winning Strategies). An *attacker strategy* is a (partial) function mapping play fragments ending at attacker positions to next positions to move to, $s_a: G^*G_a \rightarrow G$, where $g_a \rightsquigarrow f_a(\rho g_a)$ must hold for all ρg_a where s_a is defined.

A play $g_0g_1\dots \in G^\infty$ is consistent with an attacker strategy s_a if, for all its prefixes $g_0\dots g_i$ ending in $g_i \in G_a$, $g_{i+1} = s_a(g_0\dots g_i)$.

Defender strategies are defined analogously, $s_d: G^*G_d \rightarrow G$.

If s ensures a player to win, s is called a *winning strategy* for this player. The player with a winning strategy for $\mathcal{G}(g_0)$ is said to *win* $\mathcal{G}(g_0)$.

Usually, we will focus on positional strategies, that is, strategies that only depend on the most recent position, which we will type $s_a: G_a \rightarrow G$ (or $s_d: G_d \rightarrow G$, respectively).

We call the positions where a player has a winning strategy their winning region.

Definition 2.18 (Winning Region). The set $\text{Win}_a \subseteq G$ of all positions g where the attacker wins $\mathcal{G}(g)$ is called the *attacker winning region* of \mathcal{G} . The defender winning region Win_d is defined analogously.

Example 2.12 (A Simple Choice). Inspect the game in Figure 2.7, where round nodes represent defender positions and rectangular ones attacker positions. Its valid plays starting from (1) are (1), (1)[2a], (1)[2b], and (1)[2a](3). The defender can win from (1) with a strategy moving to [2b] where the attacker is stuck. Moving to [2a] instead would get the defender stuck. So, the defender winning region is $\text{Win}_d = \{(1), [2b]\}$ and the attacker wins $\text{Win}_a = \{[2a], (3)\}$.

The games we consider are positionally determined. This means, for each possible initial position, exactly one of the two players has a positional winning strategy s .

Proposition 2.6 (Determinacy). *Reachability games are positionally determined, that is, for any game, each game position has exactly one winner: $G = \text{Win}_a \cup \text{Win}_d$ and $\text{Win}_a \cap \text{Win}_d = \emptyset$, and they can win using a positional strategy.*⁴⁴

We care about reachability games because they are versatile in characterizing formal relations. Everyday inductive (and coinductive) definitions can easily be encoded as games as in the following example.




Figure 2.7: A simple game.

⁴⁴ This is just an instantiation of positional determinacy of parity games (Zielonka, 1998). Reachability games are the subclass of parity games only colored by 0.

Example 2.13 (The \leq -Game). Consider the following recursive definition of a less-or-equal operator on natural numbers in some functional programming language. (Consider `nat` to be defined as recursive data type `nat = 0 | Succ nat`.)

```
( 0 m ) = True
(Succ n 0 ) = False
(Succ n Succ m) = (n m)
```

We can think of this recursion as a game where the attacker tries to prove that the left number is bigger than the right by always decrementing the left number and challenging the defender to do the same for the right stack too. Whoever hits zero first, loses.

This means we distribute the roles such that the defender wins for output `True` and the attacker for output `False`. The two base cases need to be dead ends for one of the players.

Formally, the game \mathcal{G}_{leq} consists of attacker positions $[n, m]$ and defender positions (n, m) for all $n, m \in \mathbb{N}$, connected by chains of moves:

$$\begin{aligned} [n+1, m] &\rightsquigarrow_{\text{leq}} (n, m) \\ (n, m+1) &\rightsquigarrow_{\text{leq}} [n, m]. \end{aligned}$$

\mathcal{G}_{leq} now characterizes \leq on \mathbb{N} in the sense that: The defender wins \mathcal{G}_{leq} from $[n, m]$ precisely if $n \leq m$. (Proof: Induct over n with arbitrary m .)

The game is *boring* because the players do not ever have any choices. They just count down their way through the natural numbers till they hit $[0, m-n]$ if $n \leq m$, or $(n-m, 0)$ otherwise.

\mathcal{G}_{leq} is quite archetypical for the preorder and equivalence games we will use in the following pages. But do not worry, the following games will demand the players to make choices.

2.4.2 The Semantic Game of HML

As first bigger example of how recursive definitions can be turned into games, let us quickly look at a standard way of characterizing the semantics of HML (Definition 2.11) through a game. The defender wins precisely if the game starts from a state and formula such that the state satisfies the formula.

Definition 2.19 (HML Game). For a transition system $\mathcal{S} = (\mathcal{P}, \text{Act}, \rightarrow)$, the HML game $\mathcal{G}_{\text{HML}}^{\mathcal{S}} = (G_{\text{HML}}, G_{\text{d}}, \rightsquigarrow_{\text{HML}})$ is played on $G_{\text{HML}} = \mathcal{P} \times \text{HML}$, where the defender controls observations and negated conjunctions, that is $(p, \langle \alpha \rangle \varphi) \in G_{\text{d}}$ and $(p, \neg \bigwedge_{i \in I} \varphi_i) \in G_{\text{d}}$ (for all φ, p, I), and the attacker controls the rest.

- The defender can perform the moves:

$$\begin{aligned} (p, \langle \alpha \rangle \varphi) &\rightsquigarrow_{\text{HML}} (p', \varphi) && \text{if } p \xrightarrow{\alpha} p' \text{ and} \\ (p, \neg \bigwedge_{i \in I} \varphi_i) &\rightsquigarrow_{\text{HML}} (p, \neg \varphi_i) && \text{with } i \in I; \end{aligned}$$

- and the attacker can move:

$$\begin{aligned} (p, \neg\langle\alpha\rangle\varphi) &\rightsquigarrow_{\text{HML}} (p', \neg\varphi) && \text{if } p \xrightarrow{\alpha} p' && \text{and} \\ (p, \bigwedge_{i \in I} \varphi_i) &\rightsquigarrow_{\text{HML}} (p, \varphi_i) && \text{with } i \in I && \text{and} \\ (p, \neg\neg\varphi) &\rightsquigarrow_{\text{HML}} (p, \varphi). \end{aligned}$$

The intuition is that disjunctive constructs ($\langle\cdot\rangle \dots, \neg\bigwedge \dots$) make it easier for a formula to be true and thus are controlled by the defender who may choose which of the ways to show truth is most convenient. At conjunctive constructs ($\neg\langle\cdot\rangle \dots, \bigwedge \dots$) the attacker chooses the option that is the easiest to disprove.

Example 2.14. The game of Example 2.12 is exactly the HML game $\mathcal{G}_{\text{HML}}^{\text{SPQ}}$ for formula $\langle\tau\rangle\neg\langle a\rangle\top$ and state P of Example 2.10 with (1) := $(P, \langle\tau\rangle\neg\langle a\rangle\top)$, [2a] := $(p_a, \neg\langle a\rangle\top)$, [2b] := $(p_b, \neg\langle a\rangle\top)$, and (3) := $(p_1, \neg\top)$.

The defender winning region $\text{Win}_d = \{(P, \langle\tau\rangle\neg\langle a\rangle\top), (p_b, \neg\langle a\rangle\top)\}$ corresponds to the facts that $P \in \llbracket \langle\tau\rangle\neg\langle a\rangle\top \rrbracket$ and $p_b \in \llbracket \neg\langle a\rangle\top \rrbracket$.

As the technicalities are tangential to this thesis, we state the characterization result without proof:⁴⁵

Lemma 2.9. *The defender wins $\mathcal{G}_{\text{HML}}^S((p, \varphi))$ precisely if $p \in \llbracket \varphi \rrbracket$.*

2.4.3 The Bisimulation Game

We now can add the bottom layer of Figure 2.1: That bisimilarity can be characterized through a game. This approach has been popularized by Stirling (1996).

Definition 2.20 (Bisimulation Game). For a transition system \mathcal{S} , the *bisimulation game* \mathcal{G}_B^S ⁴⁶ is played on attack positions $G_a^B := \mathcal{P} \times \mathcal{P}$ and defense positions $G_d^B := \text{Act} \times \mathcal{P} \times \mathcal{P}$ with the following moves:

- The attacker may *swap sides*

$$[p, q] \rightsquigarrow_B [q, p],$$

- or *challenge simulation*

$$[p, q] \rightsquigarrow_B (\alpha, p', q) \text{ if } p \xrightarrow{\alpha} p';$$

- and the defender *answers simulation challenges*

$$(\alpha, p', q) \rightsquigarrow_B [p', q'] \text{ if } q \xrightarrow{\alpha} q'.$$

A schematic depiction of the game rules can be seen in Figure 2.8. From dashed nodes, the game proceeds analogously to the initial attacker position.

Example 2.15. Consider $p_{\text{even}} \sim_B p_{\text{odd}}$ of Example 2.7. The bisimulation game on this system is given by Figure 2.9:

⁴⁵ A detailed presentation of a more general HML game, also extending to recursive HML, can be found in Wortmann et al. (2015, Chapter 3).

⁴⁶ locale [Equivalence_Games.bisim_game](#)


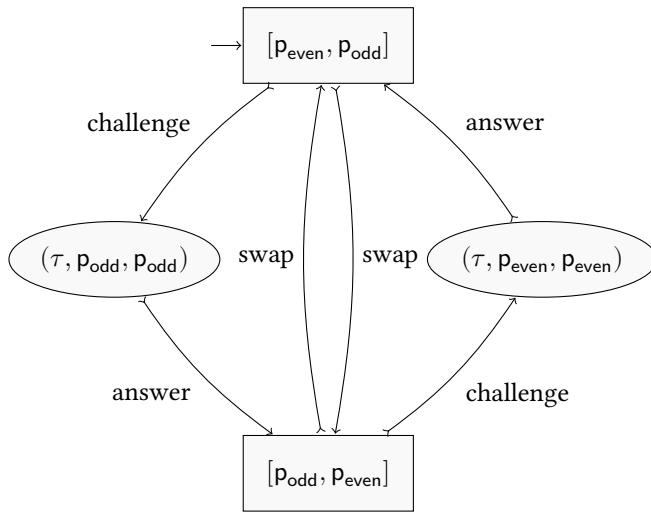


Figure 2.8: Game scheme of the bisimulation game (Definition 2.20).

Figure 2.9: Bisimulation game under $p_{\text{even}}, p_{\text{odd}}$. Moves are annotated with the game rules from which they derive.

Clearly, there is no way for the attacker to get the defender stuck. Whatever strategy the attacker chooses, the game will go on forever, leading to a win for the defender. That it is always safe for the defender to *answer* with moves to $[p_{\text{even}}, p_{\text{odd}}]$ and $[p_{\text{odd}}, p_{\text{even}}]$ corresponds to $\mathcal{R} := \{(p_{\text{even}}, p_{\text{odd}}), (p_{\text{odd}}, p_{\text{even}})\}$ being a bisimulation.

Example 2.16. Let us recall Example 2.6 of the “trolled philosophers,” where we determined Q and T to be non-bisimilar. The bisimulation game graph for the system is depicted in Figure 2.10.

The attacker can win by moving $[Q, T] \xrightarrow{B} [T, Q] \xrightarrow{B} (\tau, q_3, Q) \xrightarrow{B} [q_3, q_{AB}] \xrightarrow{B} [q_{AB}, q_3] \xrightarrow{B} (a, q_1, q_3) \not\xrightarrow{B}$. Along this sequence of positions, the defender never has a choice and is stuck in the end. The attacker exploits, that T can reach an early deadlock via $T \xrightarrow{\tau} q_3$.

Theorem 2.2 (Stirling’s Game Characterization). *The defender wins the bisimulation game \mathcal{G}_B^S starting at attacker position $[p, q]$ precisely if $p \sim_B q$.*⁴⁷

Proof. Sketch for both directions:

⁴⁷ theorem [Equivalence_Games.bisim_game.bisim_game_characterization](#)

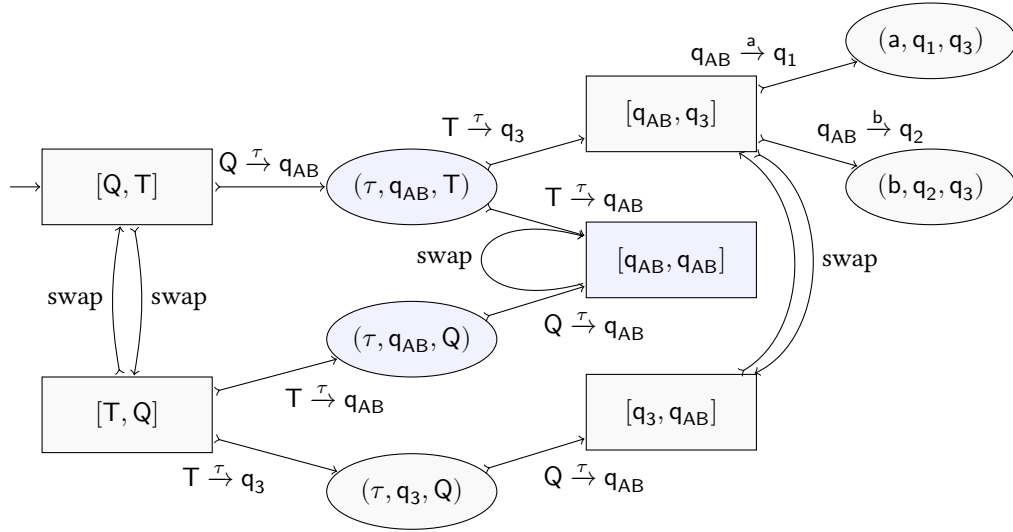


Figure 2.10: Bisimulation game on non-bisimilar states Q and T of Example 2.16. Moves are labeled by their justification. Defender-won positions are tinted blue.

- If \mathcal{R} is a symmetric simulation with $(p, q) \in \mathcal{R}$, then the following positional defender strategy is well-defined and winning from $[p, q]$:⁴⁸

$$s(\alpha, p', q) := [p', \text{choose } q'. (p', q') \in \mathcal{R} \wedge q \xrightarrow{\alpha} q'].$$

- If there is a positional defender strategy s winning from $[p, q]$, then the following relation \mathcal{R}_s with $(p_0, q_0) \in \mathcal{R}_s$ is a symmetric simulation:⁴⁹

$$\mathcal{R}_s := \{(p, q) \mid \text{there is a play } (p_0, q_0), \dots, (p, q) \text{ following strategy } s\}.$$

⁴⁸ lemma [Equivalence_Games.bisim_game.bisim_implies_defender_winning_strategy](#)

⁴⁹ lemma [Equivalence_Games.bisim_game.defender_winning_strategy_implies_bisim](#)

□

Remark 2.3. One of the big advantages of game characterizations is that they provide a way to discuss equivalence and inequivalence interactively among humans. There also are several computer game implementations of bisimulation games.

For instance, Peacock (2020) implements a game about simulation and bisimulation as well as several weaker notions. The human plays the attacker trying to point out the inequivalence of systems according to the rules of Definition 2.20. Figure 2.11 shows a screenshot of said game. It can be played on <https://www.concurrency-theory.org/rvg-game/>.




Figure 2.11: Screenshot of Peacock's bisimulation computer game.

Example 2.17. If one plays the bisimulation game of Definition 2.20 without the swapping moves, it will characterize the simulation preorder.

Consider the family of processes N_n with $n \in \mathbb{N}$. Define $N_0 := \mathbf{0}$ and $N_{n+1} := a.N_n$. Then, the simulation game played from $[N_n, N_m]$ is isomorphic to the \leq -game \mathcal{G}_{leq} from Example 2.13.

In this sense, comparisons of programs and of numbers are ... comparable.

2.4.4 How Bisimulation Game and HML Are Connected

Let us pause to think how bisimulation game and Hennessy–Milner logic connect.

You might have wondered why we even need a dedicated bisimulation game. The Hennessy–Milner theorem implies that we could directly use the HML game of Definition 2.19 to decide bisimilarity:

Definition 2.21 (Naive Bisimulation Game). We extend the Definition 2.19 by the following prefix:

1. To challenge $[p, q]$, the attacker picks a formula $\varphi \in \text{HML}$ (claiming that φ distinguishes the states) and yields to the defender (φ, p, q) .
2. The defender decides where to start the HML game:
 1. Either at $(p, \neg\varphi)$ (claiming φ to be non-distinguishing because it is untrue for p)
 2. or at (q, φ) (claiming φ to be non-distinguishing because it is true for q).
3. After that the turns proceed as prescribed by Definition 2.19.

This naive game, too, has the property that the defender wins from $[p, q]$ iff $p \sim_B q$. The downside of the game is that the attacker has infinitely many options $\varphi \in \text{HML}$ to pick from!

The proper bisimulation game of Definition 2.20, on the other hand, is *finite for finite transition systems*. Therefore, it induces decision procedures.

We will now argue that the *bisimulation game actually is a variant of the naive game*, where (1) the attacker names their formula *gradually*, and (2) the formulas stem from $\mathcal{O}_{[B]} \subseteq \text{HML}$ of Example 2.11. To this end, we will show that attacker's winning strategies imply distinguishing formulas, and that a distinguishing formula from $\mathcal{O}_{[B]}$ certifies the existence of winning attacker strategies.

Example 2.18. Let us illustrate how to derive distinguishing formulas using the game of Example 2.16.

Recall that the attacker wins by moving $[Q, T] \xrightarrow{a} [T, Q] \xrightarrow{b} (\tau, q_3, Q) \xrightarrow{c} [q_3, q_{AB}] \xrightarrow{d} [q_{AB}, q_3] \xrightarrow{e} (a, q_1, q_3) \not\xrightarrow{f}$. In Figure 2.12, we label the game nodes with the (sub-)formulas this strategy corresponds to. Swap moves become negations, and simulation moves become observations with a conjunction of formulas for each defender option. This attacker strategy can thus be expressed by $\neg\langle\tau\rangle \wedge \{\neg\langle a \rangle T\} \in \mathcal{O}_{[B]}$.

More generally, the following lemma explains the construction of distinguishing formulas from attacker strategies:

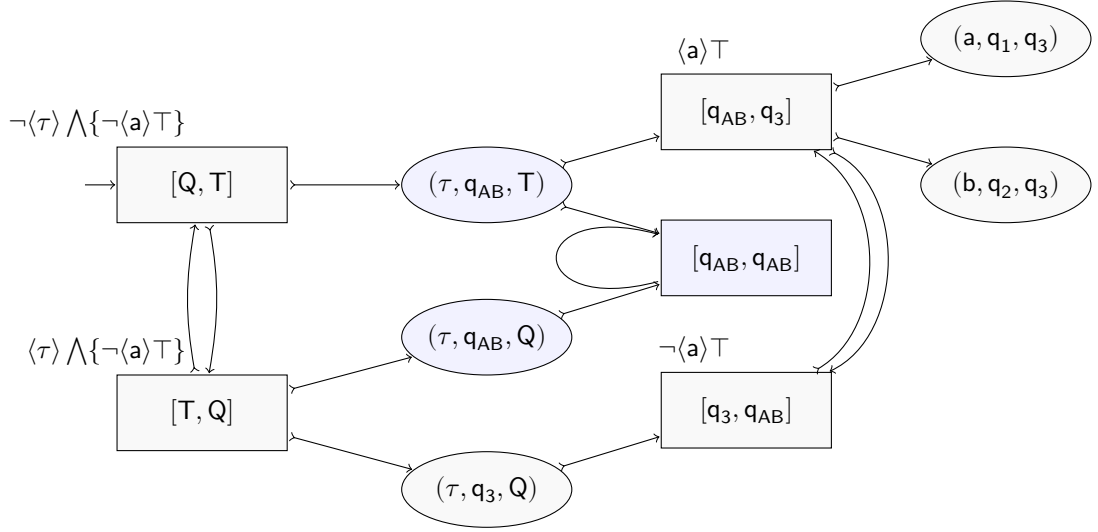


Figure 2.12: The bisimulation game of Example 2.18 with attacker formulas.

Lemma 2.10. *Let s be a positional winning strategy for the attacker on \mathcal{G}_B from $[p, q]$. Construct formulas recursively from game positions, $\varphi_s(g)$, as follows:*

$$\varphi_s([p, q]) := \begin{cases} \neg\varphi_s([q, p]) & \text{if } s([p, q]) = [q, p] \\ \langle\alpha\rangle \wedge \{\varphi_s([p', q']) \mid q \xrightarrow{\alpha} q'\} & \text{if } s([p, q]) = (\alpha, p', q) \end{cases}$$

Then $\varphi_s([p, q])$ is well-defined and distinguishes p from q . Also, $\varphi_s([p, q]) \in \mathcal{O}_{[B]}$.

Proof.

1. The given construction is well-defined as s must induce a well-founded order on game positions in order to be attacker-winning.
2. The distinction can be derived by induction on the construction of $\varphi_s([p, q])$. □

Lemma 2.11. *If $\varphi \in \mathcal{O}_{[B]}$ distinguishes p from q , then the attacker wins from $[p, q]$.*

Proof. By induction on the derivation of $\varphi \in \mathcal{O}_{[B]}$ according to the definition from Example 2.11 with arbitrary p and q .

- Case $\varphi = \langle\alpha\rangle \bigwedge_{i \in I} \varphi_i$. As φ distinguishes p from q , there must be a p' such that $p \xrightarrow{\alpha} p'$ and that $\bigwedge_{i \in I} \varphi_i$ distinguishes p' from every $q' \in \text{Der}(q, \alpha)$. That is, for each $q' \in \text{Der}(q, \alpha)$, at least one $\varphi_i \in \mathcal{O}_{[B]}$ must be false. By induction hypothesis, the attacker thus wins each $[p', q']$. As these attacker positions encompass all successors of (α, p', q) , the attacker also wins this defender position and can win from $[p, q]$ by moving there with a simulation challenge.

- Case $\varphi = \neg\varphi'$. As φ distinguishes p from q , φ' distinguishes q from p . By induction hypothesis, the attacker wins $[q, p]$. So they can also win $[p, q]$ by performing a swap. \square

Lemma 2.10 and Lemma 2.11, together with the fact that $\mathcal{O}_{|B|}$ and HML are equally distinctive (Lemma 2.8), yield:

Theorem 2.3. *The attacker wins \mathcal{G}_B from $[p, q]$ precisely if there is $\varphi \in \text{HML}$ distinguishing p from q .*

Of course, we already knew this! Theorem 2.3 is just another way of gluing together the Hennessy–Milner theorem on bisimulation (Theorem 2.1) and the Stirling’s bisimulation game characterization (Theorem 2.2), modulo the determinacy of games. We thus have added the last arrow on the right side of Figure 2.1.

2.4.5 Deciding Reachability Games

All we need to turn a game characterization into a decision procedure, is a way to decide which player wins a position. Algorithm 2.1 describes how to compute who wins a finite reachability game for each position in time linear to the size of the game graph $|\succrightarrow|$.

```

def compute_winning_region( $\mathcal{G} = (G, G_d, \succrightarrow)$ ):
  defender_options := [ $g \mapsto n \mid g \in G \wedge n = |\{g' \mid g \succrightarrow g'\}|$ ]
  attacker_win :=  $\emptyset$ 
  todo :=  $\{g \in G_d \mid \text{defender\_options}[g] = 0\}$ 
  while todo  $\neq \emptyset$ :
    g := some todo
    todo := todo  $\setminus \{g\}$ 
    if  $g \notin \text{attacker\_win}$ :
      attacker_win := attacker_win  $\cup \{g\}$ 
      for  $g_p \in (\cdot \succrightarrow g)$ :
        defender_options[ $g_p$ ] := defender_options[ $g_p$ ]  $- 1$ 
        if  $g_p \in G_a \vee \text{defender\_options}[g_p] = 0$ :
          todo := todo  $\cup \{g_p\}$ 
  Wina := attacker_win
  return Wina

```

Algorithm 2.1: Algorithm for deciding the attacker winning region Win_a of a reachability game \mathcal{G} in linear time of $|\succrightarrow|$ and linear space of $|G|$.

Intuitively, `compute_winning_region` first assumes that the defender were to win everywhere and that each outgoing move of every position might be a winning option for the defender. Over time, every position that is determined to be lost by the defender is added to a `todo` list.⁵⁰

⁵⁰ Variants of this algorithm and explanation have also been used in Bisping (2018) and Bisping et al. (2022).

At first, the defender loses immediately exactly at the defender’s dead-ends. Each defender-lost position is added to the `attacker_win` set. To trigger the recursion, each predecessor is noted as defender-lost, if it is controlled by the attacker, or the amount of outgoing defender options is decremented if the predecessor is defender-controlled. If the count of a predecessor position hits 0, the defender also loses from there.

Once we run out of `todo` positions, we know that the attacker has winning strategies exactly for each position we have visited.

The following Table 2.1 lists how Algorithm 2.1 computes the winning region $\text{Win}_a = \{[2a], (3)\}$ of the game of Example 2.11.

Table 2.1: Solving the game of Example 2.11

<code>g</code>	<code>defender_options</code>	<code>todo</code>
-	(1) \mapsto 2, (3) \mapsto 0	(3)
(3)	(1) \mapsto 2, (3) \mapsto 0	[2a]
[2a]	(1) \mapsto 1, (3) \mapsto 0	\emptyset

The inner loop of Algorithm 2.1 clearly can run at most $|\mapsto|$ -many times. Using sufficiently clever data structures, the algorithm thus shows:

Proposition 2.7. *Given a finite reachability game $\mathcal{G} = (G, G_d, \mapsto)$, the attacker’s winning region Win_a (and Win_d as well) can be computed in $O(|\mapsto|)$ time and $O(|G|)$ space.*

2.5 Discussion

This chapter has taken a tour of characterizations for standard notions of behavioral preorder and equivalence on transition systems such as trace equivalence and bisimilarity.

In particular, in constructing the relationships of Figure 2.1 for bisimilarity, we have collected the theoretical ingredients for a *certifying algorithm* to check bisimilarity of states.

Figure 2.13 describes how to not only answer the question whether two states p and q are bisimilar, but how to also either provide a bisimulation relation or a distinguishing formula for the two as certificates for the answer. (The arrows stand for computations, the boxes for data.)

In this view, the game of possible distinctions and their preventability between attacker and defender serves as the “ground truth” about the bisimilarity of two states. Bisimulation relations and modal formulas only appear as witnesses of defender and attacker strategies, mediating between game and transition system. The Hennessy–Milner theorem emerges on this level as a shadow of the determinacy of the bisimulation game (Note 1). This whole framework draws heavily from Stirling (1996).

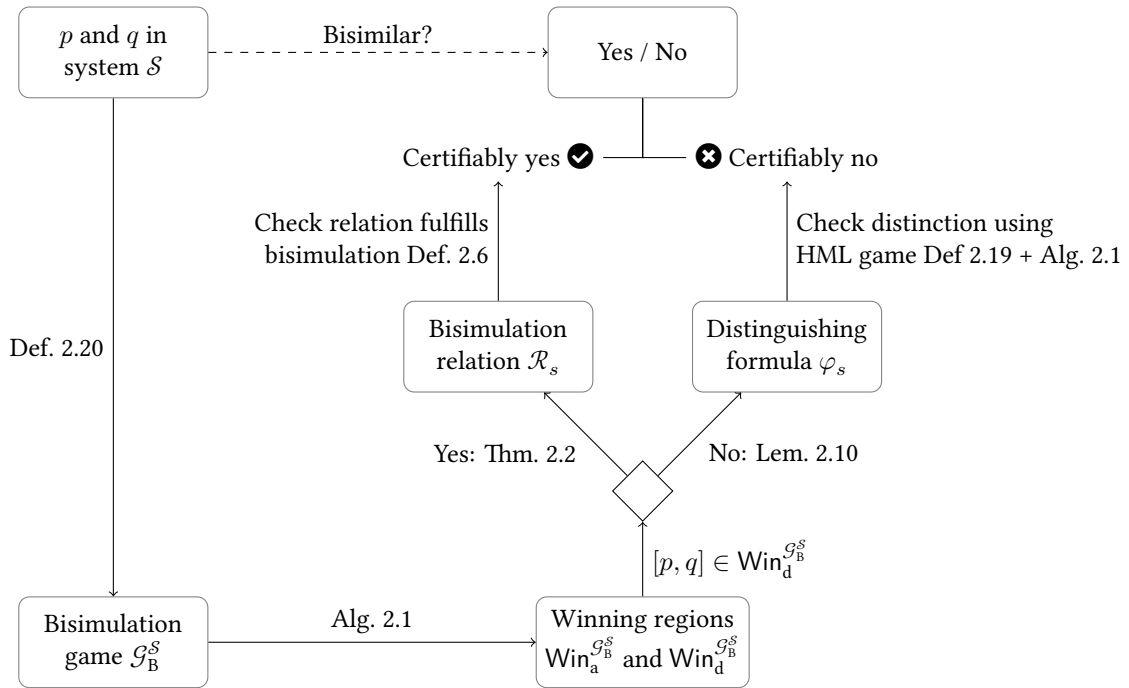


Figure 2.13: Checking bisimilarity and providing certificates.

We have disregarded the topic of axiomatic characterizations for behavioral equivalences. In these, sets of equations on process algebra terms (such as CCS) define behavioral congruences. Thus, they tend to be coupled to specific process languages and lack the genericity we pursue in starting out from transition systems.

We have observed that behavioral equivalences can be arranged in hierarchies and that these hierarchies could be handled nicely using modal characterizations to rank distinguishing powers (Subsection 2.3.3). We have also seen how to link game rules to productions in a language of potentially distinguishing formulas (Subsection 2.4.4). This departs from common textbook presentations that motivate the bisimulation game through the relational (coinductive) characterization (e.g. Sangiorgi, 2012). In later chapters, we will rather derive equivalence games from grammars of modal formulas (Note 2), to generalize the framework of Figure 2.13 for whole spectra of equivalences.

But first, we have to make formal what we mean by *equivalence spectra*.

3 Context: The Spectrum of Equivalences

We will now take a deeper dive into how groups of behavioral equivalences and preorders can be ranked in equivalence spectra.

For the following chapters, we will focus on the main “linear-time–branching-time spectrum“ for the semantics of concrete processes treated by Glabbeek (1990), the so-called “strong spectrum.” However, we will order the equivalences using the approach of parameterized *notions of observability* from the later Glabbeek (1993).

The first core idea of this chapter is:

i Idea 3: Notions of observability bring order

Groups of equivalences can be defined and ranked in lattices of notions of observability.

In particular, Section 3.2 will introduce such a spectrum, forming a hierarchy of modal logics for the *strong spectrum*.

From there, we will quickly run into the second core idea:

i Idea 4: We have a Spectroscopy problem

We can ask what equivalences from a spectrum are the most fitting to relate two states.

Section 3.3 will define the problem formally and give lower-bounds for its complexity on the strong spectrum.

3.1 Observability Hierarchies

Let us begin to pick up the idea from Subsection 2.3.3 that modal logics can nicely rank equivalences. The intuition is that HML sublogics capture what we consider to be *observable*. The more we mark as observable, the finer the resulting equivalence relations do become.

3.1.1 Understanding the Equivalence Hierarchy through Modal Logics

As promised, we revisit the hierarchy between bisimilarity, similarity, and trace equivalence of Subsection 2.2.3, modally. So far, we have only looked into the characterization of bisimilarity through the whole of HML in Theorem 2.1 or through $\mathcal{O}_{[B]}$ in Example 2.11.

Definition 3.1. We define the two HML-sublogics \mathcal{O}_T , the linear positive fragment, and \mathcal{O}_S , the positive fragment, by the grammars starting at φ^T and φ^S .

$$\begin{aligned}\varphi^T & ::= \langle \alpha \rangle \varphi^T \mid \top \\ \varphi^S & ::= \langle \alpha \rangle \varphi^S \mid \bigwedge_{i \in I} \langle \alpha_i \rangle \varphi_i^S\end{aligned}$$

The logics characterize trace and simulation preorder (and equivalence):

Lemma 3.1. $p \preceq_T q$ precisely if $p \preceq_{\mathcal{O}_T} q$.⁵¹

Lemma 3.2. $p \preceq_S q$ precisely if $p \preceq_{\mathcal{O}_S} q$.⁵²

Clearly, $\mathcal{O}_T \subset \mathcal{O}_S \subset \text{HML}$. So, Proposition 2.4 that sublogics will equate at least as much as their parent logic, yields another way of establishing the entailment hierarchy between bisimilarity, similarity, and trace equivalence of Subsection 2.2.3.⁵³

While the relational definitions for (bi-)similarity of Definition 2.7 and the denotational definition of trace equivalence Definition 2.5 live in different worlds, the two equivalences become naturally comparable in the modal realm.

The modal view also reveals an intuitive hierarchy of “testing scenarios” for the equivalences, framed as black box tests:

Trace equivalence matches an observer that can see *sequences of events*. They *just watch repeated executions* of the program, but are oblivious to possibilities and decisions.

Similarity matters to an experimenter that can also explore different *branches of possibilities*. This is valid if the experimenter can somehow *copy the system state* during the execution.

Bisimilarity captures that the experimenter can moreover determine if a future course of events is *impossible at some point*. This means that the experimenter can not only copy the execution state but also exhaustively test *every possibility* of how the system may continue.

But such levels of observability do not need to be linear, as we will see in the next subsection ...

3.1.2 Incomparabilities

A well-known and natural notion of equivalence is that of *failure equivalence*. Intuitively, a failure says that the experimenter may follow a trace and see

⁵¹ theorem [HML_Spectrum.lts.observations_traces_characterizes_trace_preorder](#)

⁵² theorem [HML_Spectrum.lts.observations_simulation_characterize_simulation_preorder](#)

⁵³ Taking the two previous lemmas together with Theorem 2.1 for \sim_B and HML.

which actions are *impossible* at its end. Its standard definition is based on *failure* denotations:

Definition 3.2 (Failures). The set of failures of a process $\text{Failures}(p)$ is recursively defined as

- $((), X) \in \text{Failures}(p)$ if $X \cap \text{Ini}(p) = \emptyset$,
- $(\alpha \cdot \bar{w}, X) \in \text{Failures}(p)$ if there is p' with $p \xrightarrow{\alpha} p'$ and $(\bar{w}, X) \in \text{Failures}(p')$.

For instance, the failure $(\tau, \{a\})$ is in $\text{Failures}(P)$ but not in $\text{Failures}(Q)$ on Figure 2.3.

But would it not be nice if we could prevent the invention of new mathematical objects as denotations for every new notions of observability we consider? Fortunately, we can save the work, by directly employing modal logics:

Definition 3.3. We define failure observations \mathcal{O}_F by the grammar:

$$\varphi^F ::= \langle \alpha \rangle \varphi^F \quad | \quad \bigwedge_{i \in I} \neg \langle \alpha_i \rangle \top$$

Clearly, this encompasses what we may observe via traces, but is something that we cannot observe using simulation. We consider \preceq_F given by $\preceq_{\mathcal{O}_F}$.

The distinguishing failure $(\tau, \{a\})$ can be as $\langle \tau \rangle \wedge \{ \neg \langle a \rangle \top \} \in \mathcal{O}_F$ in HML. The formula distinguishes P from Q on Figure 2.3, but $P \preceq_S Q$ (cf. Example 2.5). On the other hand, no failure from \mathcal{O}_F , distinguishes Q from P , but $Q \not\preceq_S P$.

As a consequence, simulation preorder and failure preorder are *incomparable*, that is, neither one implies the other. The same is true of the corresponding equivalences: similarity and failure equivalence. The situation is summed up by the non-linear hierarchy in Figure 3.1. After a quick glance at the diamond-like figure, it probably comes as no surprise, what kind of mathematical structure we employ to handle the hierarchy: Lattices.

3.1.3 Lattices

To handle non-linearity, we will be working with lattices of notions of behavioral equivalence. The following definition gives the preliminaries to talk about this kind of partial orders.

Definition 3.4 (Lattices, Bounds, Chains). A *lattice* is a partially ordered set (B, \leq) , where there are greatest lower bounds $\inf\{b_1, b_2\}$ and least upper bounds $\sup\{b_1, b_2\}$ between each pair of elements $b_1, b_2 \in B$.

The *greatest lower bound* of a set $B' \subseteq B$ is called its *infimum*, $\inf B'$. It refers to the greatest element $b \in B$ such that $b \leq b'$ for all $b' \in B'$.⁵⁴

Dually, the *least upper bound* of a set $B' \subseteq B$ is called its *supremum*, $\sup B'$. It refers to the least element $b \in B$ such that $b \geq b'$ for all $b' \in B'$.

For the pair-wise infimum we also use infix notation $b_1 \sqcap b_2 = \inf\{b_1, b_2\}$, and analogously $b_1 \sqcup b_2 = \sup\{b_1, b_2\}$.




Figure 3.1: Preorder/equivalence hierarchy.

⁵⁴ Note that not necessarily $b \in B'$!

If a lattice (B, \leq) not only allows infima and suprema for pairs but for any set $B' \subseteq B$, it is called *complete*. We say \sqcap -complete or \sqcup -complete if only one of the two is true.

We call totally ordered subset $B' \subseteq B$ a *chain*.

Dually, some $B' \subseteq B$ with no two elements $b_1, b_2 \in B'$ such that $b_1 \leq b_2$ is called an *anti-chain*.

Example 3.1 (Subset Lattice). Given any set B , its subsets and ordered by set inclusion, $(2^B, \subseteq)$, form a complete lattice.

The *greatest lower bound* is given by set intersection $\bigcap B'$ with $B' \subseteq B$. The *least upper bound* is set union $\bigcup B'$ with $B' \subseteq B$. $\emptyset \in 2^B$ is the least element, $B \in 2^B$ is itself the greatest bound.

Consider the subset lattice over $B = \{1, 2, 3\}$. It is “cube-like,” as can be seen in its Hasse diagram in Figure 3.2. An example of a (maximal) chain would be $\{\emptyset, \{1\}, \{1, 2\}, B\}$ (the nodes connected by a blue dotted line in the figure), because its members are ordered linearly $\emptyset \subset \{1\} \subset \{1, 2\} \subset B$. $\{\{1\}, \{2\}, \{3\}\}$ forms a (maximal) anti-chain (the nodes connected by a red dashed line in the figure), because its members do not include each other. Their respective subsets are chains / anti-chains as well.




Figure 3.2: Lattice of subsets from Example 3.1.

Example 3.2 (Vector Lattice). Given linearly-ordered set (B, \leq_B) , its n -ary Cartesian product with pointwise order (B^n, \leq) forms a lattice, where $b \leq b'$ iff $b_k \leq b'_k$ for all $k \in \{1, \dots, n\}$. Greatest lower bounds and least upper bounds can be transferred pointwisely from B .

For instance, pairs of natural numbers, (\mathbb{N}^2, \leq) , form a lattice, as visualized in Figure 3.3. It is \sqcap -complete, that is, for any set from \mathbb{N}^2 , one can pick a greatest lower bound. However, the lattice is not \sqcup -complete: For instance, the set $\mathbb{N} \times \{0\}$ has no upper bound. If we take the natural numbers extended with an upper bound ∞ , \mathbb{N}_∞ , as basis, then $(\mathbb{N}_\infty^2, \leq)$ forms a complete lattice.




Figure 3.3: Visualization of the infinitary grid of \mathbb{N}^2 (and in gray, \mathbb{N}_∞^2) in Example 3.2.

3.2 The Linear-Time–Branching-Time Spectrum

Van Glabbeek’s two papers on the “linear-time–branching-time spectrum” (1993, 2001) show how all common equivalences can be understood to form a lattice of sublanguages of HML (and a variant of HML for equivalences with silent steps). His hierarchy of equivalences derives from a hierarchy of *notions of observability*.⁵⁵ We will introduce a similar construction: at first, in a generic form, then for the strong spectrum of Glabbeek (2001).

3.2.1 Spectra as Observability Lattices

As we will discuss various equivalence spectra (van Glabbeek (1993, 2001) already gives two different ones), let us first introduce an abstract description of such spectra.

⁵⁵ In particular, the weak spectrum (Glabbeek, 1993) makes this really formal.

Definition 3.5 (Equivalence Spectrum). An *equivalence spectrum* $(\mathbf{N}, \leq, \mathcal{O}_{N \in \mathbf{N}})$ consists of

- a lattice of notions of observability, \mathbf{N} , partially ordered by $\leq \subseteq \mathbf{N} \times \mathbf{N}$, and
- corresponding logics $\mathcal{O}_N : 2^{\text{HML}}$ for $N \in \mathbf{N}$.

$\mathcal{O}_{N \in \text{notions}}$ must be monotonic, that is: for any two notions $N, M \in \mathbf{N}$, it holds that

$$N \leq M \text{ implies } \mathcal{O}_N \subseteq \mathcal{O}_M.$$

Let us use our new definition to construct a subset lattice as in Example 3.1 to recreate the hierarchy of Figure 3.1.

Example 3.3. Consider the notions

$$\mathbf{N}^{\text{simple}} := 2^{\{(\wedge), (\neg)\}},$$

ordered by subset inclusion, and the family of observation languages $\mathcal{O}_{N \in \mathbf{N}^{\text{simple}}}$ given by the family of grammars with some conditional productions:

$$\begin{array}{l} \varphi^N ::= \top \\ \quad | \langle \alpha \rangle \varphi^N \\ \quad | \bigwedge_{i \in I} \varphi_i^N \quad \text{if } (\wedge) \in N \\ \quad | \bigwedge_{i \in I} \neg \langle \alpha_i \rangle \top \quad \text{if } (\neg) \in N \\ \quad | \neg \varphi^N \quad \text{if } \{(\wedge), (\neg)\} = N. \end{array}$$

Clearly, $N \subseteq M$ implies $\mathcal{O}_N \subseteq \mathcal{O}_M$ for $N, M \in \mathbf{N}^{\text{simple}}$. We obtain the diamond hierarchy of Figure 3.4. It matches the diamond of Figure 3.1, but this time, the hierarchy is an effect of the ordered notions.

While the incomparable languages of Subsection 3.1.2 form no lattice, e.g. $\mathcal{O}_S \cup \mathcal{O}_F \neq \text{HML}$, the *notions* of the present Example 3.3 do form a lattice, as $\{(\wedge)\} \cup \{(\neg)\} = \{(\wedge), (\neg)\}$. This is one of the reasons why it is convenient to add *notions of observability* as an abstraction layer.

We can also ask what the least notion is to include a specific formula:

Definition 3.6 (Syntactic Expressiveness Price). In the context of a \sqcap -complete spectrum $(\mathbf{N}, \leq, \mathcal{O}_{N \in \mathbf{N}})$, the *syntactic expressiveness price* of a formula φ that appears in one of the logics is defined as

$$\text{expr}(\varphi) := \inf\{N \in \mathbf{N} \mid \varphi \in \mathcal{O}_N\}.$$

Thinking of the lattice of notions as a hierarchy of how difficult it is to tell processes apart, we consider this as a kind of “price tag” to put on formulas depending on their syntactic complexity. Higher syntactic complexity allows formula sets of higher expressiveness.

In this view, a trace formula is *cheaper* than a failure formula. Using Example 3.3: $\text{expr}(\langle \tau \rangle \langle a \rangle) = \emptyset \subset \{(\neg)\} = \text{expr}(\langle \tau \rangle \wedge \{ \neg \langle a \rangle \})$, which captures

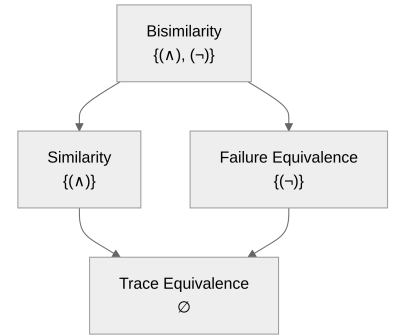


Figure 3.4: Hierarchy of simple notions of equivalence.

that we need a strictly smaller part of the grammar to construct the trace formula.

3.2.2 The Strong Notions of Observability

Now we can approach the *strong spectrum*. We will encode its notions as a \mathbb{N}_∞ -vector lattice as in Example 3.2. To cover all common behavioral preorders, we use six dimensions, counting certain syntactic features of HML-formulas:

1. Formula modal depth of observations.
2. Formula nesting depth of conjunctions.
3. Maximal modal depth of deepest positive clauses in conjunctions.
4. Maximal modal depth of the other positive clauses in conjunctions.
5. Maximal modal depth of negative clauses in conjunctions.
6. Formula nesting depth of negations.

Definition 3.7 (Strong Spectrum). We define the *strong notions of observability* using vectors of extended naturals

$$\mathbf{N}^{\text{strong}} := \mathbb{N}_\infty^6,$$

ordered by pointwise comparison of vector components, and the family of strong observation languages $\mathcal{O}_{N \in \mathbf{N}^{\text{strong}}}^{\text{strong}}$ given by the parameterized grammar starting from φ^N :⁵⁶

⁵⁶ primrec [Priced_HML.formula_of_price](#)

$$\begin{aligned} \varphi^N & ::= \top \\ & | \langle \alpha \rangle \varphi^{N-\hat{e}_1} \\ & | \bigwedge \{ \psi^{N-\hat{e}_2}, \psi^{(N-\hat{e}_2) \sqcap (\infty, \infty, N_4, \infty, \infty, \infty)}, \psi^{(N-\hat{e}_2) \sqcap (\infty, \infty, N_4, \infty, \infty, \infty)}, \dots \} \\ \psi^N & ::= \langle \alpha \rangle \varphi^{(N \sqcap (N_3, \infty, \infty, \infty, \infty, \infty))-\hat{e}_1} \\ & | \neg \langle \alpha \rangle \varphi^{(N \sqcap (N_5, \infty, \infty, \infty, \infty, \infty))-\hat{e}_1-\hat{e}_6} \end{aligned}$$

The productions only exist if the respective recursive invocations are defined on the domain of notions. For instance, $\varphi^N \rightsquigarrow \langle \alpha \rangle \varphi^{N-\hat{e}_1}$ is no valid production for $N = (0, 1, 0, 0, 0, 0)$ because $(-1, 1, 0, 0, 0, 0) \notin \mathbf{N}^{\text{strong}}$.

Example 3.4. The smallest notion to cover the failure formula of Subsection 3.1.2 would be $(2, 1, 0, 0, 1, 1)$, that is,

$$\langle \tau \rangle \wedge \{ \neg \langle a \rangle \top \} \in \mathcal{O}_{(2,1,0,0,1,1)^{\text{strong}}}.$$

This is because the formula has two levels of modal observations, where the inner one is negated. The negation is wrapped in a conjunction with no positive clauses. A visualization for how $\text{exp}^{\text{strong}}(\langle \tau \rangle \wedge \{ \neg \langle a \rangle \top \}) = (2, 1, 0, 0, 1, 1)$ (according to Definition 3.6) comes together is given in the upcoming Fig-

ure 3.5. Formally, the reason is that the following derivation is optimal:

$$\begin{aligned}
\varphi^{(2,1,0,0,1,1)} &\rightsquigarrow \langle \tau \rangle \varphi^{(1,1,0,0,1,1)} \\
&\rightsquigarrow \langle \tau \rangle \wedge \{ \psi^{(1,0,0,0,1,1)} \} \\
&\rightsquigarrow \langle \tau \rangle \wedge \{ \neg \langle a \rangle \varphi^{(0,0,0,0,1,0)} \} \\
&\rightsquigarrow \langle \tau \rangle \wedge \{ \neg \langle a \rangle \top \}.
\end{aligned}$$

If we reexamine the grammar of \mathcal{O}_F in Definition 3.3, we notice that the formulas it can produce match those of $\mathcal{O}_{(\infty,1,0,0,1,1)}^{\text{strong}}$ of the strong spectrum in Definition 3.7.




Figure 3.5: Pricing formula $\langle \tau \rangle \wedge \{ \neg \langle a \rangle \top \}$ with syntactic expressiveness $e = (2, 1, 0, 0, 1, 1)$.

An example for the pricing of a more complex tree-like formula is given in Figure 3.6.




Figure 3.6: Pricing formula $\langle \tau \rangle \wedge \{ \langle ec_A \rangle \langle lc_A \rangle \top, \langle \tau \rangle \top, \neg \langle ec_B \rangle \top \}$ with syntactic expressiveness $e = (3, 1, 2, 1, 1, 1)$.

Lemma 3.3. *The strong spectrum of Definition 3.7 contains the notions of behavioral equivalence we have discussed so far.*

1. The observation language $\mathcal{O}_{(\infty,0,0,0,0,0)}^{\text{strong}}$ exactly matches the characterization of traces \mathcal{O}_T from Definition 3.1 and thus characterizes trace pre-order.⁵⁷
2. The observation language $\mathcal{O}_{(\infty,\infty,\infty,\infty,0,0)}^{\text{strong}}$ exactly matches the characterization of simulation observations \mathcal{O}_S from Definition 3.1 and thus characterizes simulation.⁵⁸
3. The observation language $\mathcal{O}_{(\infty,\infty,\infty,\infty,\infty,\infty)}^{\text{strong}}$ matches HML in distinctiveness and thus characterizes bisimilarity.

⁵⁷ theorem [Priced_Spectrum.lts.traces_priced_characterization](#)

⁵⁸ theorem [Priced_Spectrum.lts.simulation_priced_characterization](#)

4. The observation language $\mathcal{O}_{(\infty,1,0,0,1,1)}^{\text{strong}}$ exactly matches failure observations \mathcal{O}_{F} of Definition 3.3.

Proof. The only non-trivial case is (3) for bisimilarity:

Observe that $\mathcal{O}_{[\text{B}]} \subseteq \mathcal{O}_{(\infty,\infty,\infty,\infty,\infty,\infty)}^{\text{strong}}$ by examining its grammar in Example 2.11. As $\mathcal{O}_{[\text{B}]}$ already has complete HML-distinctiveness by Lemma 2.8, so must its superlogic $\mathcal{O}_{(\infty,\infty,\infty,\infty,\infty,\infty)}^{\text{strong}}$. \square

So far, we have only established that the six-dimensional spectrum also covers the notions that Example 3.3 has already covered—in a more complicated way. The extra dimensions will pay off in the next subsection.

3.2.3 The Strong Linear-Time–Branching-Time Spectrum

Using the six dimensions of Definition 3.7, we can assign coordinates to all other common notions of the strong linear-time–branching-time spectrum.

Definition 3.8 (Strong Linear-Time–Branching-Time Spectrum). Coordinates with respect to the notions of Definition 3.7 for the common notions of behavioral equivalence and preorder in the strong linear-time–branching-time spectrum are given in Figure 3.7.⁵⁹

For the rest of the thesis, we will take the equivalences as defined by the coordinates as canonical. Lemma 3.3 has established that the coordinates of traces, simulation, failures, and bisimulation match the definitions one commonly finds in the literature. To establish the same for the other notions, would mean a lot of (repetitive) definitions, which mostly reproduce van Glabbeek’s survey (2001). Mattes (2024) proves in Isabelle/HOL that the coordinate system of the conference version of this spectrum (Bisping, 2023a) match distinctiveness of the modal characterizations in (Glabbeek, 2001).

⁵⁹ When writing vectors in labels and figures, we omit the parentheses (...) for readability.

Warning

I might add such a marathon if there’s strong popular demand.

There are a few standard questions that come to mind for people who are familiar with the various spectra of equivalences when seeing Figure 3.7. The following remarks address these points.

Remark 3.1 (Selection of Notions). At the core, we treat the same notions as Glabbeek (2001). But we feature a slightly more modern selection.

Our spectrum additionally includes strong versions of *impossible futures* (Voorhoeve & Mauw (2001)) and *revivals* (Roscoe (2009)) as equivalences whose relevance has only been noted after the publication of Glabbeek (2001).

On the other hand, we glimpse over completed trace, completed simulation, and possible worlds observations like Kučera & Esparza (1999), who studied properties of “good” observation languages. These notions would need a

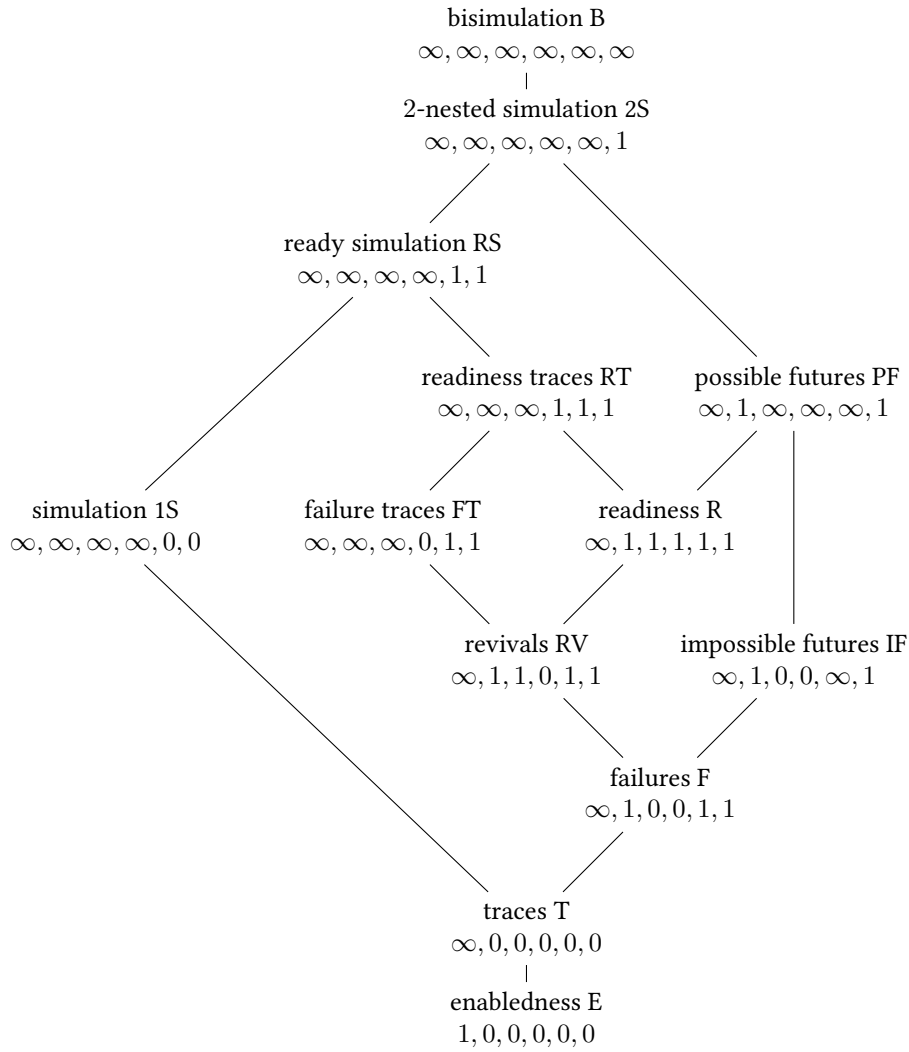


Figure 3.7: Hierarchy of common equivalences/preorders ordered by observability coordinates.

different HML grammar, featuring exhaustive conjunctions $\bigwedge_{a \in Act} \varphi_a$, where the φ_a are deactivated actions for completed traces, and more complex trees for possible worlds.

Remark 3.2 (Ambiguous Coordinates). For many of the logics in Figure 3.7, there are multiple coordinates that characterize the same logic. For instance, due to the second dimension (conjunctions) being set to 0 for traces T , the higher dimensions do not matter and any coordinate $N = (\infty, 0, N_3, N_4, N_5, N_6)$ will lead to the same observation language $\mathcal{O}_N^{\text{strong}} = \mathcal{O}_T^{\text{strong}}$.

Indeed, Figure 3.7 always selects the least coordinate to characterize a sublogic, which ensures that domination of coordinates in the figure and entailment between behavioral preorders coincide.

Remark 3.3 (Other Coordinates). We have singled out a handful of coordinates. Many other coordinates will still correspond to distinct equivalences. For instance, we could consider $N^{2T} = (2, 0, 0, 0, 0, 0)$, preordering states that cannot be distinguished by traces up to a length of 2. But it is difficult to make a case for such a “notion of equivalence,” which washes away differences of future behavior.

Two kinds of depth-bounded families, however, are common in the literature, and can also be placed in our spectrum:

- k -step bisimilarity: $(k, \infty, \infty, \infty, \infty, \infty)$ is a k -observation approximation of bisimilarity that sometimes appears in proofs.⁶⁰
- k -nested similarity: $(\infty, \infty, \infty, \infty, \infty, k - 1)$ for $k > 1$ defines a spectrum of modal quantifier alternation depth between similarity and bisimilarity.

⁶⁰ TODO: look up example

Remark 3.4 ((In-)finitary Variants). One can introduce more dimensions to the spectrum with respect to the possibility of infinitary observations. Our choice focuses on natural and most common versions of the equivalences, in particular: similarity and bisimilarity with unbounded (infinitary) branching and trace-like notions with finitary depth. Notions in Figure 3.7 correspond precisely to those without superscripts in the infinitary linear-time–branching-time spectrum of Glabbeek (2001, Figure 9).

3.2.4 Non-Intersectionality

The strong spectrum of Definition 3.8 is much richer than the diamond spectrum from Example 3.3. Still, its observation languages form no lattice. For instance, the lines of simulation and failures join at ready simulation—and their coordinates as well $(\infty, \infty, \infty, \infty, 0, 0) \sqcup (\infty, 1, 0, 0, 1, 1) = (\infty, \infty, \infty, \infty, 1, 1)$. But $\mathcal{O}_S^{\text{strong}} \cup \mathcal{O}_F^{\text{strong}} \neq \mathcal{O}_{RS}^{\text{strong}}$ and this makes a difference:

Example 3.5. Consider the CCS processes $a.(a.b + a)$ and $a.a.b + a.a$. They cannot be told apart by $\mathcal{O}_S^{\text{strong}}$ or $\mathcal{O}_F^{\text{strong}}$ and thus are simulation and failure equivalent (and moreover even ready-trace equivalent).

Still, the formula $\langle a \rangle \wedge \{ \langle a \rangle \wedge \{ \neg \langle b \rangle \}, \langle a \rangle \langle b \rangle \} \in \mathcal{O}_{RS}^{\text{strong}}$ distinguishes the first process from the second. Therefore, the processes are not ready-simulation equivalent.

What this shows is that one cannot prove two states to be ready-equivalent by showing that they are equated by simulation and failures:

$$\sim_S \cap \sim_F \not\subseteq \sim_{RS}.$$

The relationship between the characterized equivalences is *non-intersectional*.

In general, multiple preorders may relate two states without this entailing a stronger equivalence. So the question “Which equivalence from a spectrum relates two states?” is to simple, one has to ask: “Which equivalences relate two processes?”

This motivates the *spectroscopy problem*.

3.3 Spectroscopy

Now that we have a formal way of describing equivalence spectra, we can make formal the *spectroscopy problem*—the core topic of this thesis. We will also collect first thoughts on its complexity.

3.3.1 The Spectroscopy Problem

The problem has originally been introduced in Bisping et al. (2022) as the “abstract observation preorder problem” with respect to modal characterizations of the strong spectrum. We here reintroduce it in a more generic form.

Definition 3.9 (Spectroscopy Problem). In the context of a transition system \mathcal{S} and a spectrum $(\mathbb{N}, \leq, \mathcal{O}_{N \in \mathbb{N}})$, the *spectroscopy problem* asks:

Input States p and q .

Output Set of notions $N_{p,q} \subseteq \mathbb{N}$, such that $p \preceq_{\mathcal{O}_N} q$ for each $N \in N_{p,q}$.

Example 3.6. For the “trolled philosophers” of Example 2.6, we have determined that the systems are simulation-preordered, but not bisimilar, that is, $Q \preceq_S T$, but $Q \not\preceq_B T$. The first fact implies $Q \preceq_T T$.

But what about other notions from the strong spectrum of Subsection 3.2.3? Besides similarity, there might well be incomparable or finer notions that too preorder Q to T !

The solution to the spectroscopy problem on Q and T is $N_{Q,T}^{\text{strong}} = \{N \in \mathbb{N}^{\text{strong}} \mid N \not\preceq (2, 2, 0, 0, 2, 2)\}$. A minimal formula to distinguish Q from T with this coordinate would be $\bigwedge \{ \neg \langle \tau \rangle \wedge \{ \neg \langle a \rangle \top \} \}$. (The following chapters

will reveal how to reliably arrive at this knowledge, in particular, the minimality.)

Because the coordinate of 2-nested simulation, $2S = (\infty, \infty, \infty, \infty, \infty, 1)$ is not less or equal to $(2, 2, 0, 0, 2, 2)$, we arrive at $Q \not\leq_{2S} T$, which implies *all* preorders of Figure 3.7 except for bisimilarity.

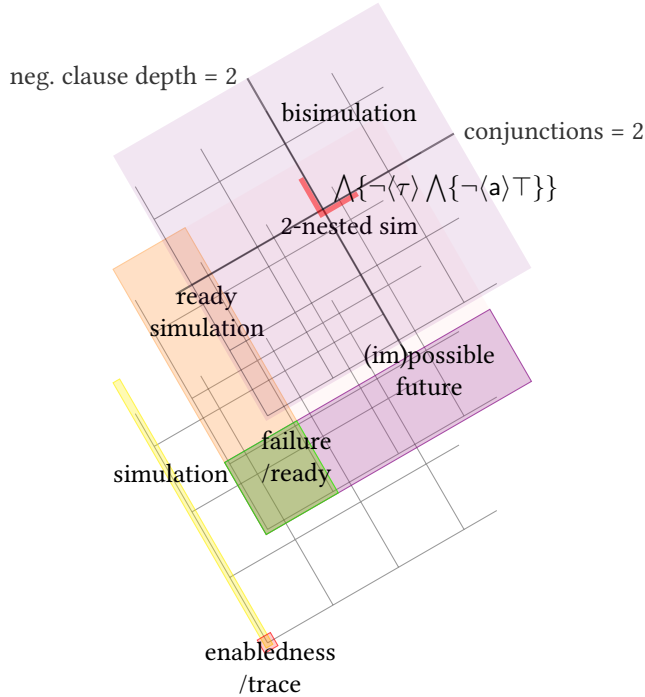


Figure 3.8: Cross section of the strong spectrum showing the dimensions conjunctions, negative clause depth, and negation depth. All preorders below the red mark at $(2, 2, 2)$ hold in Example 3.6.

Note that we have expressed $N_{Q,T}^{\text{strong}}$ through a negation (“... $\not\leq (2, 2, 0, 0, 2, 2)$ ”). The reason is that a positive description is usually unwieldy. In this (comparably easy) case, we could for example list the half-spaces below the cheapest distinction, and this would read: $N_{Q,T}^{\text{strong}} = (\{0, 1\} \times \mathbb{N}_{\infty}^5) \cup (\mathbb{N}_{\infty} \times \{0, 1\} \times \mathbb{N}_{\infty}^4) \cup (\mathbb{N}_{\infty}^4 \times \{0, 1\} \times \mathbb{N}_{\infty}) \cup (\mathbb{N}_{\infty}^5 \times \{0, 1\})$.

Technically, it is convenient to not compute $N_{p,q}$ directly. Rather we aim to construct the *Pareto front* of minimal notions that do not hold, $\text{Min}(N \setminus N_{p,q})$. The Pareto front serves as a unique representation, from which $N_{p,q}$ can be constructed as complement of the upwards closure $N \setminus \uparrow \text{Min}(N \setminus N_{p,q})$ (where $\uparrow B := \{b \mid \exists b' \in B. b' \leq b\}$ is the upwards closure for B subset of some partially-ordered set). Clearly, Pareto fronts form *anti-chains* and appear naturally in optimization problems.

All spectra we are concerned with are well-quasi ordered, which means that each $\text{Min}(N \setminus N_{p,q})$ must be finite in size (Kruskal, 1972) and thus “more

handy” than the full sets $N_{p,q}$ or $N \setminus N_{p,q}$.

So, effectively, we will be asking: What are the minimal notions to distinguish p from q —and then often talk about the converse: The most-fitting notions to preorder or equate the states. Everything else is implied.

3.3.2 Spectroscopy as Abstract Subtraction

Another way of viewing the spectroscopy problem is that we aim to compute an abstracted kind of *difference* between programs.

Definition 3.10 (Observations and Difference). On a transition system \mathcal{S} , the possible observations of a state, $\llbracket \cdot \rrbracket^{\bullet} : \mathcal{P} \rightarrow 2^{\text{HML}}$, are defined as:

$$\llbracket p \rrbracket^{\bullet} := \{\varphi \in \text{HML} \mid p \in \llbracket \varphi \rrbracket\}.$$

The *difference* between p and q is defined as:

$$\Delta(p, q) := \llbracket p \rrbracket^{\bullet} \setminus \llbracket q \rrbracket^{\bullet}.$$

$\Delta(p, q)$ expresses the set of observations one could make of p that one cannot make of q . This set will be empty, when the states are bisimilar, or infinite, otherwise.

With this notation, we could rephrase Definition 2.12:

Proposition 3.1. *Two states p and q are preordered with respect to a sublogic $\mathcal{O} \subseteq \text{HML}$:*

$$p \preceq_{\mathcal{O}} q \iff \Delta(p, q) \cap \mathcal{O} = \emptyset.$$

A spectroscopy then is about computing some abstraction Δ_{α} such that $N \in \Delta_{\alpha}(p, q)$ precisely if $\Delta(p, q) \cap \mathcal{O}_N \neq \emptyset$.

3.3.3 Complexities

What complexities to expect when deciding spectroscopy problems on finite systems? Details depend, of course, on the specific spectrum and flavor of Hennessy–Milner logic we are concerned with. Still, solving the spectroscopy problem cannot be easier than solving the covered individual equivalence problems.

To get a first idea, let us examine the complexities of common equivalence checking problems in the strong spectrum. The rule of thumb is that trace-like equivalences are PSPACE-complete and bisimilarities are P-complete (Bal-cázar et al., 1992; Hüttel & Shukla, 1996).

Bisimilarity finds itself in a *valley* of tractability, if we look at a cross section through the equivalence spectrum as in Figure 3.9. The best known bisimilarity algorithms for finite-state transition systems take $O(|\rightarrow| \log |S|)$ time (usually deriving from Paige & Tarjan, 1987).

For coarser simulation-like equivalences, the best known algorithms need $O(|\rightarrow| |S|)$ time (Ranzato & Tapparo, 2010).⁶¹

⁶¹ Or $O(|\rightarrow| |S|_{\sim_s})$ to name the bound as Ranzato & Tapparo (2010) present it.




Figure 3.9: Bisimilarity's complexity valley.

The finer graph-isomorphism equivalence (Definition 2.8) again is harder with the best known solution (Babai, 2016) in quasi-polynomial time $2^{O((\log n)^3)}$.

There however are little strict hardness results at this level of granularity. So, better time complexities for graph isomorphism, bisimilarity, and similarity are conceivable (albeit improbable). Groote et al. (2023) show that at least partition-refinement algorithms for bisimilarity cannot do better than $O(|\rightarrow| \log |S|)$. In a recent preprint, Groote & Martens (2024) establish that similarity is strictly more complex than bisimilarity.

The trivial equivalences at the end of the cross section, identity ID and universal equivalence U, can be solved directly. Enabledness equivalence E can as well be computed quite quickly by just comparing outgoing transitions.

In this thesis, we solve the spectroscopy problem for the strong and weak spectrum. So, we must be at least as complex as the equivalences between bisimilarity and universal equivalence, boxed in Figure 3.9. Consequently, the spectroscopy problem for the standard equivalence spectra is PSPACE-hard.

3.4 Discussion

In this chapter, we have formalized how to handle *spectra of equivalence* (Note 3), and instantiated the approach to the strong spectrum of Glabbeek (2001). From there, we have introduced the *spectroscopy problem*, which asks for notions to preorder compared states (Note 4).

By formulating the problem in terms of a lattice over \mathbb{N}_∞ -vectors, the family of qualitative strong preorder/equivalence problems becomes a single quantitative problem: The spectroscopy problem for the strong spectrum.

We have already laid the groundwork to shift the semantic question of equivalence into a syntactic question of the shape of distinguishing formulas.

The core ideas of this section have already been explored in Bisping et al. (2022) and Bisping (2023a) for the strong spectrum. However, in my prior work the expressiveness prices played a more crucial role. Here, we instead opted for a parameterized grammar to define notions and their observations \mathcal{O}_N . In this grammar, we force observations in clauses of conjunctions

and count \top as part of the $\mathbf{0}$ -notion. These are mostly superficial changes to streamline the following presentation. We have shown that traces, failures, simulation and bisimulation equivalence as defined by the notion coordinates match their textbook definitions.

So far, we have a problem and confidence that its solution conveys information about equivalences of the strong spectrum. Subsection 3.3.3 has established that spectroscopy complexity must be at least PSPACE on the strong spectrum.

However, there is a polynomial-time part of spectrum around bisimilarity and (ready) similarity. In the next two chapters, we will first solve the spectroscopy problem for this polynomial slice and then extend to the whole strong spectrum. After that, we will also consider the weak spectrum.

4 Approach: Equivalence Problems as Energy Games

Time to get real. This chapter will demonstrate my core approach by solving the spectroscopy problem in an *easy* instance.

It will be easy in two regards: Firstly, complexity-wise—we narrow our focus on equivalences of polynomial-time. Secondly, conceptually—we just use the bisimulation game of Chapter 2, almost directly!

We already know that winning attacker strategies in the bisimulation game correspond to distinguishing formulas, and that these can be computed quite straightforwardly. All we need on top of this is a way of quantifying the amount of syntactic expressiveness in these formulas *during the game*. For this, we employ the first core idea of this chapter:

i Idea 5: Quantitative games for quantitative problems

Spectra of equivalence problems can be encoded as energy games.

Energy games are games in which players have limited resources that can be used up or recharged during moves. Players running out of a resource lose the game. After introducing such games in Section 4.1, we will prove in Section 4.2 that several core equivalences can be characterized as defender's winning energy budgets by adding three-dimensional energies to the bisimulation game.

The second core idea is how to compute these winning budgets:

i Idea 6: Computing cheapest wins

Attacker's winning budgets in energy reachability games can be computed by a generalized shortest paths algorithm.

Section 4.3 will provide an algorithm to solve a range of energy-game-like quantitative problems as long as the energy updates can be *undone* through a *Galois connection*, which is a generalization of invertibility on monotonic functions.

By combining the two contributions, we arrive at a polynomial-time solu-

tion of the spectroscopy problem for the polynomial slice of the strong spectrum.

Effectively, we adapt the game framework of Figure 2.13 to not treat one equivalence, but a spectrum of equivalences. This approach is summarized in Figure 4.1.

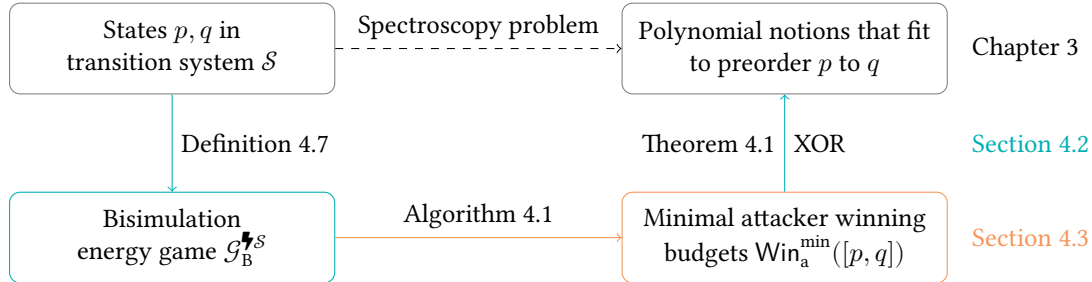


Figure 4.1: How we will employ energy games to solve the spectroscopy problem on an easy spectrum.

4.1 Energy Games

Energy games extend games as discussed in Section 2.4 with resources of the players, called energies. We will focus on reachability games with an energy-bounded attacker. In most publications, energy games update vector-valued energy levels by vector addition or subtraction. We work with more general *monotonic energy games* (Subsection 4.1.1) and then zoom in on *declining energy games* (Subsection 4.3), which include updates that combine vector components by taking their minimum. The latter are exactly what we need for the following Section 4.2.

4.1.1 Monotonic Energy Games

To introduce energy games, we generalize the definitions on games (Definition 2.15 and the following) to include *energy levels*. We define how winning regions then become quantitative winning *budgets*, which can still be characterized inductively.

Definition 4.1 (Energy Reachability Game). Given a partially ordered set of energies (\mathcal{E}, \leq) , an *energy reachability game* \mathcal{G}^\sharp is a reachability game $(G, G_d, \rightsquigarrow)$ extended by an edge labeling of energy updates $upd: (\rightsquigarrow) \rightarrow (\mathcal{E} \rightarrow (\{\perp\} \cup \mathcal{E}))$.

We demand the following for monotonic energy games:

- $\perp \notin \mathcal{E}$ is final in the sense that $upd(m)(\perp) = \perp$.
- All update functions $u = upd(m)$ with $m \in \rightsquigarrow$ are monotonic and upward-closed with respect to \leq . More formally, this means that for

any energies $e, e' \in \mathcal{E}$, if $e \in \text{dom}(u)$ and $e \leq e'$, then $e' \in \text{dom}(u)$ and $u(e) \leq u(e')$, where $\text{dom}(u) := \{e \in \mathcal{E} \mid u(e) \neq \perp\}$.

We denote by $\mathcal{G}^\blacktriangleright(g_0, e_0)$ the game played from starting position $g_0 \in G$ with initial energy level $e_0 \in \mathcal{E}$.

Definition 4.2 (Energy Level). For a finite play $\rho = g_0 g_1 \dots g_{n-1} \in G^*$ of $\mathcal{G}^\blacktriangleright(g_0, e_0)$, its energy level $\text{EL}(\rho)$ is computed recursively:

- $\text{EL}(g_0) := e_0$
- $\text{EL}(g_0 \dots g_{i+1}) := \text{upd}(g_i, g_{i+1})(\text{EL}(g_0 \dots g_i))$

For infinite plays $\rho \in G^\omega$, we define energy levels to equal \perp .

We consider the attacker to be energy-bounded and understand $\text{EL}(\rho) = \perp$ to mean that they have run out of energy. Thus, we declare plays with $\text{EL}(\rho) = \perp$ to be won by the defender (even if they are stuck).

Strategies and winning strategies work exactly as in the energy-less scenario. Additionally, we lift positional strategies of Definition 2.17 to be *energy-positional* in the sense that they pick next moves depending on the current energy level, i.e. $f: G_a \times \mathcal{E} \rightarrow G$.

Instead of a winning region as in Definition 2.18, we define *winning budgets* for game positions:

Definition 4.3 (Winning Budgets). For each position $g_0 \in G$ of energy game $\mathcal{G}^\blacktriangleright$, the *attacker winning energy budgets*, $\text{Win}_a(g_0) \subseteq \mathcal{E}$ are the energies e_0 where the attacker wins $\mathcal{G}^\blacktriangleright(g_0, e_0)$. The defender winning budgets Win_d are defined analogously.

Classical energy games use \mathbb{N} -vectors for energies and vector addition for updates as in the following example:

Example 4.1. Consider vector energies $\mathcal{E} = \mathbb{N}^2$ with point-wise order and the energy game with graph in Figure 4.2. Edges are labeled by update vectors $\vec{u} \in \mathbb{Z}^2$, each representing an update function

$$e \mapsto \begin{cases} e + \vec{u} & \text{if } e + \vec{u} \geq \mathbf{0} \\ \perp & \text{otherwise.} \end{cases}$$

Can the attacker win from g_1 with energy $(0, 0)$?

No, as all outgoing paths would lead to \perp -energy and thus to the defender winning.

But if the attacker starts with budget $(0, 2)$, they can take the upper path to g_4 with energy $(2, 1)$, from where both defender options lead to the defender being stuck in g_6 .

Also, if the attacker starts with $(2, 1)$, they win through the lower g_3 -path.

Moving to g_4 directly, would be more expensive than the g_2/g_3 -alternatives—so we can disregard this option.




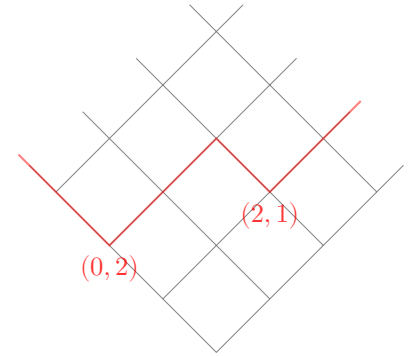
Figure 4.2: Simple energy game of Example 4.1.

Starting with less energy means that the defender has an option of bankrupting the attacker at g_4 .

Summing up our observations, $(0, 2)$ and $(2, 1)$ define a Pareto front of minimal budgets where the attacker wins from g_4 , as depicted in Figure 4.3.

We can notice two things from the example:

1. A purely positional strategy would not suffice for the defender to win reliably at g_4 . Say the game starts with energy $(1, 1)$, that is, with a budget where the defender should win. Then, the level with either be $(0, 1)$ or $(3, 0)$ at g_4 . The defender has to either move to g_5 or g_6 , depending on the play so far. So, either the defender must know the history, or at least the current energy level to make an ideal decision.
2. Attacker's winning budgets are upward-closed, and defender budgets are downward-closed.

Figure 4.3: Attacker's Pareto front for g_1 of example energy game.

Both are effects of Definition 4.1. More formally, the winning budgets are characterized by the following two propositions:

Proposition 4.1. *Given an energy game \mathcal{G} , the attacker winning budgets $\text{Win}_a^{\mathcal{G}}$ are point-wise upward-closed, that is, $\uparrow \text{Win}_a^{\mathcal{G}}(g) = \text{Win}_a^{\mathcal{G}}(g)$ for all positions $g \in G$.⁶²*

Proposition 4.2. *Given an energy game \mathcal{G} , the attacker winning budgets $\text{Win}_a^{\mathcal{G}}$ are characterized inductively by the following rules:⁶³*

$$\frac{g_a \in G_a \quad g_a \rightsquigarrow g' \quad \text{upd}(g_a, g')(e) \in \text{Win}_a^{\mathcal{G}}(g')}{e \in \text{Win}_a^{\mathcal{G}}(g_a)}$$

$$\frac{g_d \in G_d \quad \forall u, g'. \quad g_d \rightsquigarrow g' \longrightarrow \text{upd}(g_d, g')(e) \in \text{Win}_a^{\mathcal{G}}(g')}{e \in \text{Win}_a^{\mathcal{G}}(g_d)}$$

The inductive definition is nice as it allows for local proves of how the attacker wins. Defender's wins on the other hand would dually be a coinductive concept because the defender wins loops.⁶⁴

⁶² This is formalized as Lemma 3.2 in Lemke (2024).

⁶³ Proved as Theorem 3 in Lemke (2024).

⁶⁴ By the way, to those wondering: It is intentional that this thesis aims to achieve all its constructions without recurring to coinduction. Therefore, coinduction is only mentioned in side remarks and is not required to understand what is going on.

Remark 4.1. Ordinary reachability games can be seen as a special case of energy games with trivial energies, e.g. $\mathcal{E} = \{1\}$ and $\text{upd}(\cdot) = \text{id}_{\mathcal{E}}$.

Remark 4.2. We deviate in two important points from other work on energy games:

- In the literature, it is more common to consider the defender instead of the attacker to be energy-bounded (e.g. [Fahrenberg et al., 2011](#)). This choice follows the intuition that there is a party with scarce resources that wants to keep the system running. For our purposes, however we want to bound the resources of an attacker, which is no fundamental change. Still, one cannot have both at the same time: [Kupferman & Shamash Halevy \(2022\)](#), studying both-bounded energy games, establish that their winner is only decidable if both parties have only one-dimensional energies.
- We define energy-reachability games more abstractly than is common: In other publications, energy games work on numbers with component-wise updates as seen in [Example 4.1](#). We do not demand this as, in the next sections, we will need updates that combine information from different components of the energy vector. Even prior generalizations on energy update functions such as [Ésik et al. \(2013\)](#) do not allow sufficient flexibility in this regard.

4.1.2 Declining Energy Games

We now turn to special kind of monotonic energy games with vector-valued energy levels where updates will never increase energy levels in any component. An N -dimensional *declining energy game* is played on energy vectors:

Definition 4.4 (Energies and Energy Updates). For dimensionality N , the set of *energies*, \mathbf{En} , is given by $(\mathbb{N} \cup \{\infty\})^N$.

The set of *energy update labels*, \mathbf{Up} , contains $(u_1, \dots, u_N) \in \mathbf{Up}$ where each component u_k is a symbol of the form

- $u_k \in \{-1, 0\}$ (relative update), or
- $u_k = \min_D$ where $D \subseteq \{1, \dots, N\}$ and $k \in D$ (minimum selection update).

Applying an update to an energy, $\text{upd}(u, e)$, where $e = (e_1, \dots, e_N) \in \mathbf{En}$ and $u = (u_1, \dots, u_N) \in \mathbf{Up}$, yields a new energy vector e' where k th components $e'_k := e_k + u_k$ for $u_k \in \mathbb{Z}$ and $e'_k := \min_{d \in D} e_d$ for $u_k = \min_D$. Updates that would cause any component to become negative yield \perp .

Example 4.2. Consider the update label $(\min_{\{1,2\}}, 0, -1)$.

- $\text{upd}((\min_{\{1,2\}}, 0, -1), (1, 1, 0))$ equals \perp because of the last component.
- $\text{upd}((\min_{\{1,2\}}, 0, -1), (2, 1, 1))$ equals $(1, 1, 0)$.

- $\text{upd}((\min_{\{1,2\}}, 0, -1), (1, 1, 1))$ equals $(1, 1, 0)$, as well.
- There is no $e \in \text{En}$ such that $\text{upd}((\min_{\{1,2\}}, 0, -1), e) = (1, 0, 0)$, because the second component would demand a lower first component.

In summary, $\text{upd}((\min_{\{1,2\}}, 0, -1), \cdot)$ is neither injective nor surjective with respect to En . The same is true for most other u and $\text{upd}(u, \cdot)$ with a min -update component.

Definition 4.5 (Declining Energy Game). Given a weight labeling $w: (\succrightarrow) \rightarrow \text{Up}$, an N -dimensional *declining energy game* \mathcal{G}^\bullet is an energy reachability game with energies $\mathcal{E} := \text{En}$ and with edges labeled $\text{upd}(m) := \text{upd}(w(m), \cdot)$.

We write $g \xrightarrow{u} g'$ for a move $g \succrightarrow g'$ labeled by weight $w(g, g') = u$.

Declining energy games differ from more common energy games with vector addition updates as seen in Example 4.1 due to the non-invertibility of updates we encountered in Example 4.2.

Proposition 4.3. *In this model, energy levels can only decline.*

- Updates may only decrease energies, $\text{upd}(u, e) \leq e$.
- Updates are monotonic as demanded: If $e_A \leq e_B$, then $\text{upd}(u, e_A) \leq \text{upd}(u, e_B)$.

4.2 Characterizing the P-Time Part of the Spectrum

As hinted at in Subsection 3.3.3, parts of the strong equivalence spectrum can be decided in polynomial time, whereas others require polynomial space (and thus effectively exponential time). In this section, we will show how to adapt the standard bisimulation game of Definition 2.20 to decide not just bisimulation, but *all polynomial-time* equivalences of the strong spectrum *at once*.

4.2.1 The Polynomial Slice

First, let us make explicit which equivalences we include in the polynomial-time slice of the spectrum. (The fact that they indeed can be decided efficiently is mostly well-known, but also itself a corollary of this section.)

To define the polynomial slice, we only use three dimensions:

1. modal depth
2. depth of negative clauses, and
3. nesting depth of negations (that is, the first, fifth and sixth dimension of the strong spectrum Definition 3.7).

Components for other dimensions of the strong spectrum of the previous chapter are effectively set to ∞ .

Definition 4.6 (Polynomial-Time Strong Spectrum). We denote as *polynomial-time strong spectrum* the projection of the strong spectrum (Definition 3.7) to the first, fifth, and sixth dimension, thus by notions

$$\mathbf{N}^{\text{poly-strong}} := \mathbb{N}_{\infty}^3,$$

and the family of strong observation languages

$$\mathcal{O}_{(N_1, N_2, N_3) \in \mathbf{N}^{\text{poly-strong}}}^{\text{poly-strong}} := \mathcal{O}_{(N_1, \infty, \infty, \infty, N_2, N_3)}^{\text{strong}}.$$

In effect, $\mathcal{O}_{N \in \mathbf{N}^{\text{poly-strong}}}$ can be characterized by φ^N in the following grammar:

$$\begin{array}{l} \varphi^N ::= \top \\ \quad | \langle \alpha \rangle \varphi^{N - \hat{e}_1} \\ \quad | \bigwedge \{ \psi^N, \psi^N, \psi^N \dots \} \\ \psi^N ::= \langle \alpha \rangle \varphi^{N - \hat{e}_1} \\ \quad | \neg \langle \alpha \rangle \varphi^{(N \sqcap (N_2, \infty, \infty)) - \hat{e}_1 - \hat{e}_3} \end{array}$$

Figure 4.4 gives names to coordinates that correspond to notions discussed previously.

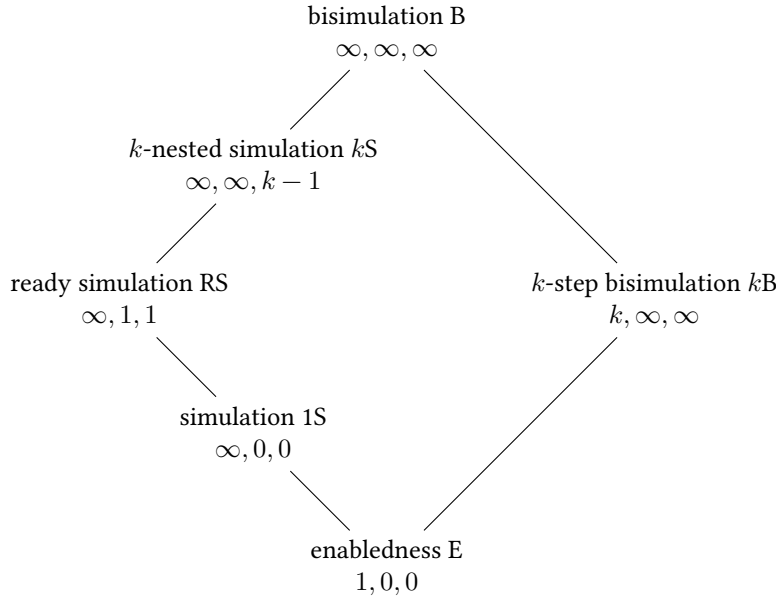


Figure 4.4: Hierarchy of polynomial-time decidable equivalences/preorders.

Except for enabledness, E, all coordinates after unfolding them for $\mathbf{N}^{\text{strong}}$ are identical to the ones used previously. When one unfolds the coordinate for enabledness $(1, 0, 0)$ to the original strong system, $(1, \infty, \infty, \infty, 0, 0)$, it differs from the one previously used, namely, $(1, 0, 0, 0, 0, 0)$. The characterized language is strictly more expressive, as it contains conjunctions of enabled

actions (e.g. $\bigwedge\{\langle a \rangle, \langle b \rangle\}$), but it is equally distinctive.⁶⁵

⁶⁵ lemma [Priced_Spectrum.lts.enabledness_conjunctions_are_neutral](#)

4.2.2 The Bisimulation Energy Game

Let us upgrade the bisimulation game of Definition 2.20 with energies. The claim is that the resulting game \mathcal{G}_B^\sharp will characterize all equivalences of the polynomial-time strong spectrum in the sense that $p \preceq_{\mathcal{O}_e} q$ for $e \in \mathbb{N}^{\text{poly-strong}}$ precisely if the defender wins \mathcal{G}_B^\sharp from $[p, q]$ with budget e .

As we have seen in Subsection 2.4.4, the game moves match the distinctive power of productions in the $\mathcal{O}_{|B|}$ -grammar of Example 2.11. Consequently, we can *count* the use of HML constructs in the game. As the game stands, we can meaningfully count the three dimensions used in Subsection 4.2.1.⁶⁶

We add the computations that happen in the grammar of Definition 4.6 as energy updates to the bisimulation game and obtain:

Definition 4.7 (Bisimulation Energy Game). For a transition system \mathcal{S} , the *bisimulation energy game* $\mathcal{G}_B^{\sharp\mathcal{S}}$ is played on the same graph as the *bisimulation game* $\mathcal{G}_B^{\mathcal{S}}$ (of Definition 2.20), but the moves are weighted by the following energy updates:

- Attacker swaps are counted as a negation and limit the further observations to the depth of negative clauses:

$$[p, q] \xrightarrow{\min_{\{1,2\}, 0, -1}^\sharp} [q, p].$$

- Simulation challenges count as an observation (using up the budget for modal depth):

$$[p, q] \xrightarrow{-1, 0, 0}^\sharp (\alpha, p', q) \quad \text{if } p \xrightarrow{\alpha} p'.$$

- And defender answers come for free:

$$(\alpha, p', q) \xrightarrow{0, 0, 0}^\sharp [p', q'] \quad \text{if } q \xrightarrow{\alpha} q'.$$

The first dimension thus bounds how often the attacker may challenge simulation down the road, the third limits how often they may swap sides, and the middle dimension bounds the amount of simulation moves after a swap.

Example 4.3. Let us label the bisimulation game of Example 2.18 (distinguishing the “trolled philosophers”) with energy updates. Figure 4.5 shows the game graph, also exhibiting the cheapest formulas with regards to the polynomial spectrum that correspond to attacker winning strategies.

The attacker wins \mathcal{G}_B^\sharp from $[Q, T]$ if they start out with an energy budget of $(2, 2, 2)$ or above. But if the budget does not dominate this bound, the attacker loses. For instance, neither $(1, \infty, \infty)$ nor $(\infty, \infty, 1)$ is enough. The second bound implies that the budgets, $(\infty, 1, 1)$ and $(\infty, 0, 0)$ are won by the defender as well.

⁶⁶ We cannot hope to count the other dimensions of Definition 3.7: the bisimulation game washes away the amount of *necessary* conjunctions as its HML has conjunctions under each observation, even though they might not be necessary for a distinguishing formula.

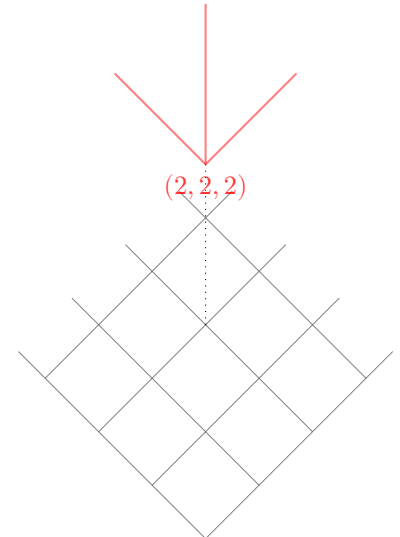


Figure 4.6: Attacker’s Pareto front for $[Q, T]$ of the example bisim energy game.




Figure 4.5: The bisimulation energy game of Example 4.3 with energy updates and implied formulas.

On the game-side, this tells us is that the attacker needs at least two simulation moves (after swaps) and two swaps to tell Q apart from T. Telling apart T from Q is slightly easier, only requiring one swap (and only one simulation move after the swap).

But, due to the correspondence of the game to modal formulas and of the energy updates to the pricing of formulas in the spectrum, the attacker winning budgets tell us more: They reveal that one needs two observations and two (or more) negations to distinguish Q from T using formulas from $\mathcal{O}_{[B]}$.

Thus, $\mathbf{N}_{Q,T}^{\text{poly-strong}} = \{N \in \mathbf{N}^{\text{poly-strong}} \mid N \not\leq (2, 2, 2)\}$ is the solution for the (polynomial-time strong) spectroscopy problem. Q is preordered to T by 2-nested simulation, 1-step bisimulation, and all notions below; but distinguished by all above, in particular, 3-nested simulation and 2-step bisimulation.

This is in line with the solution we thought of for the original strong spectrum in Example 3.6.

4.2.3 Correctness of Characterization

We now prove that the bisimulation energy game indeed characterizes the equivalences of the polynomial-time strong spectrum. In following chapters, we will prove stronger results for more general games in more detail. The approach is to generalize the connection between winning attacks and distinguishing formulas that we have already explored in Lemma 2.10 and Lemma 2.11 to energies and spectrum.

Definition 4.8 (Strategy Formulas for \mathcal{G}_B^g). The set of *strategy formulas* for a game position g and a budget e , $\text{Strat}_B(g, e)$, in the context of a bisimulation

energy game $\mathcal{G}_B^{\uparrow S}$ is defined inductively by the following rules:

$$\frac{[p, q] \xrightarrow{-1, 0, 0}_B^{\uparrow} (\alpha, p', q) \quad e' = \text{upd}((-1, 0, 0), e) \in \text{Win}_a^{\mathcal{G}_B^{\uparrow}}((\alpha, p', q)) \quad \varphi \in \text{Strat}_B((\alpha, p', q), e')}{\langle \alpha \rangle \varphi \in \text{Strat}_B([p, q], e)}$$

$$\frac{\forall q' \in \text{Der}(q, \alpha). \quad (\alpha, p', q) \xrightarrow{0, 0, 0}_B^{\uparrow} [p', q'] \wedge e \in \text{Win}_a^{\mathcal{G}_B^{\uparrow}}([p', q']) \wedge \psi_{q'} \in \text{Strat}_B([p', q'], e)}{\left(\bigwedge_{q' \in \text{Der}(q, \alpha)} \psi_{q'} \right) \in \text{Strat}_B((\alpha, p', q), e)}$$

$$\frac{[p, q] \xrightarrow{\min_{\{1, 2\}, 0, -1}_B^{\uparrow}} [q, p] \quad e' = \text{upd}((\min_{\{1, 2\}, 0, -1}), e) \in \text{Win}_a^{\mathcal{G}_B^{\uparrow}}([q, p]) \quad \langle \alpha \rangle \varphi \in \text{Strat}_B([q, p], e')}{\neg \langle \alpha \rangle \varphi \in \text{Strat}_B([p, q], e)}$$

Effectively, this definition generalizes the construction of distinguishing formulas on \mathcal{G}_B we introduced in Lemma 2.10. So the formulas we encountered before also serve as examples of such strategy formulas.

Lemma 4.1. *If $e \in \text{Win}_a^{\mathcal{G}_B^{\uparrow}}([p, q])$, then there is $\varphi \in \text{Strat}_B([p, q], e)$ with $\varphi \in \mathcal{O}_e^{\text{poly-strong}}$, $p \in \llbracket \varphi \rrbracket$ and $q \notin \llbracket \varphi \rrbracket$.*

Proof. We prove the following, more general, property by induction.

1. If $e \in \text{Win}_a^{\mathcal{G}_B^{\uparrow}}([p, q])$, then there is $\psi \in \text{Strat}_B([p, q], e)$ of form $\psi = \langle \alpha \rangle \varphi'$ or $\neg \langle \alpha \rangle \varphi'$ with $\bigwedge \{\psi\} \in \mathcal{O}_e^{\text{poly-strong}}$, $p \in \llbracket \psi \rrbracket$ and $q \notin \llbracket \psi \rrbracket$.
2. If $e \in \text{Win}_a^{\mathcal{G}_B^{\uparrow}}((\alpha, p', q))$, then there is $\varphi \in \text{Strat}_B((\alpha, p', q), e)$ of form $\varphi = \bigwedge \Psi$ with $\varphi \in \mathcal{O}_e^{\text{poly-strong}}$, $p' \in \llbracket \varphi \rrbracket$ and $q \notin \llbracket \langle \alpha \rangle \varphi \rrbracket$.

Proof by induction on the inductive characterizations of attacker winning budgets Proposition 4.2.

1. Assume $e \in \text{Win}_a^{\mathcal{G}_B^{\uparrow}}([p, q])$. This must be due to one of the following moves:
 - Case $[p, q] \xrightarrow{\uparrow}_B [q, p]$ with $e' = \text{upd}((\min_{\{1, 2\}, 0, -1}), e) \in \text{Win}_a^{\mathcal{G}_B^{\uparrow}}([q, p])$. By induction hypothesis on $[q, p]$, there is $\psi \in \text{Strat}_B([q, p], e')$ with $\bigwedge \{\psi\} \in \mathcal{O}_{e'}^{\text{poly-strong}}$, $q \in \llbracket \psi \rrbracket$ and $p \notin \llbracket \psi \rrbracket$. Consider the possible forms of ψ :
 - $\psi = \langle \alpha \rangle \varphi'$. By Definition 4.8 of Strat_B , we obtain $\neg \langle \alpha \rangle \varphi' \in \text{Strat}_B([p, q], e)$. Because of the semantics of HML, $\neg \langle \alpha \rangle \varphi'$ must distinguish p from q . Also, $\bigwedge \{\neg \langle \alpha \rangle \varphi'\} \in \mathcal{O}_e^{\text{poly-strong}}$, by the calculations of the grammar in Definition 4.6.
 - $\psi = \neg \langle \alpha \rangle \varphi'$. This can only be in $\text{Strat}_B([q, p], e')$ due to φ' being in $\text{Strat}_B([p, q], \text{upd}((\min_{\{1, 2\}, 0, -1}), e'))$. Clearly, φ' must distinguish p from q and $\bigwedge \{\varphi'\} \in \mathcal{O}_e^{\text{poly-strong}}$.
 - Case $[p, q] \xrightarrow{\uparrow}_B (\alpha, p', q)$ with $p \xrightarrow{\alpha} p'$ and $e' = \text{upd}((-1, 0, 0), e) \in \text{Win}_a^{\mathcal{G}_B^{\uparrow}}((\alpha, p', q))$. By induction hypothesis on (α, p', q) , there is $\bigwedge \Psi \in \text{Strat}_B((\alpha, p', q), e')$ and $\bigwedge \Psi \in \mathcal{O}_{e'}^{\text{poly-strong}}$ with

$p' \in \llbracket \langle \alpha \rangle \wedge \Psi \rrbracket$ and $q \notin \llbracket \langle \alpha \rangle \wedge \Psi \rrbracket$. By Definition 4.8, we obtain $\langle \alpha \rangle \wedge \Psi \in \text{Strat}_B([p, q], e)$. Due to the semantics of HML and $p \xrightarrow{\alpha} p'$, $p \in \llbracket \langle \alpha \rangle \wedge \Psi \rrbracket$, thus distinguishing p from q . Also, $\langle \alpha \rangle \wedge \Psi \in \mathcal{O}_e^{\text{poly-strong}}$ as $\langle \alpha \rangle \wedge \Psi \in \mathcal{O}_{e-\hat{e}_1}^{\text{poly-strong}}$.

2. Assume $e \in \text{Win}_a^{\mathcal{G}_B^\sharp}((\alpha, p', q))$. This means that e suffices for the attacker to win every q' -move $(\alpha, p', q) \xrightarrow{\mathcal{G}_B^\sharp} [p', q']$ with $q \xrightarrow{\alpha} q'$ that the defender might take, with $e \in \text{Win}_a^{\mathcal{G}_B^\sharp}([p', q'])$. For each q' , we employ the induction hypothesis to obtain $\psi_{q'} \in \text{Strat}_B([p', q'], e)$ of form $\psi_{q'} = \langle \alpha \rangle \varphi'$ or $\neg \langle \alpha \rangle \varphi'$ with $\bigwedge \{\psi_{q'}\} \in \mathcal{O}_e^{\text{poly-strong}}$, $p' \in \llbracket \psi_{q'} \rrbracket$ and $q' \notin \llbracket \psi_{q'} \rrbracket$. By Definition 4.8, $(\bigwedge_{q' \in \text{Der}(q, \alpha)} \psi_{q'}) \in \text{Strat}_B((\alpha, p', q), e)$. By the HML-semantics, $(\bigwedge_{q' \in \text{Der}(q, \alpha)} \psi_{q'})$ must be true for p' , and also taking the semantics of observation, $\langle \alpha \rangle (\bigwedge_{q' \in \text{Der}(q, \alpha)} \psi_{q'})$ false for q . Moreover, $(\bigwedge_{q' \in \text{Der}(q, \alpha)} \psi_{q'}) \in \mathcal{O}_e^{\text{poly-strong}}$ by the grammar. \square

Lemma 4.2. *If there is $\varphi \in \mathcal{O}_e^{\text{poly-strong}}$ with $p \in \llbracket \varphi \rrbracket$ and $q \notin \llbracket \varphi \rrbracket$, then $e \in \text{Win}_a^{\mathcal{G}_B^\sharp}([p, q])$.*

Proof. By induction on the grammar of $\mathcal{O}_e^{\text{poly-strong}}$.

Consider the cases of $\varphi \in \mathcal{O}_e^{\text{poly-strong}}$ with $p \in \llbracket \varphi \rrbracket$ and $q \notin \llbracket \varphi \rrbracket$.

- $\varphi = \top$. This cannot be the case as $q \notin \llbracket \varphi \rrbracket$ and $\llbracket \top \rrbracket = \mathcal{P}$.
- $\varphi = \langle \alpha \rangle \varphi'$ with $\varphi' \in \mathcal{O}_{e-\hat{e}_1}^{\text{poly-strong}}$. As φ distinguishes p from q , there must be a $p' \in \llbracket \varphi' \rrbracket$ such that $p \xrightarrow{\alpha} p'$ and, for all $q' \in \text{Der}(q, \alpha)$, $q' \notin \llbracket \varphi' \rrbracket$. Due to the induction hypothesis, $e - \hat{e}_1 \in \text{Win}_a^{\mathcal{G}_B^\sharp}([p', q'])$ for all such $q' \in \text{Der}(q, \alpha)$. Therefore, $e - \hat{e}_1 \in \text{Win}_a^{\mathcal{G}_B^\sharp}((\alpha, p', q))$. As the attacker can move $[p, q] \xrightarrow{\hat{e}_1} (\alpha, p', q)$, this proves $e \in \text{Win}_a^{\mathcal{G}_B^\sharp}([p, q])$ by Proposition 4.2.
- $\varphi = \bigwedge \{\Psi\}$ one of the $\psi \in \Psi$ must be false for q . Consider the possible forms of ψ :
 - Case $\psi = \langle \alpha \rangle \varphi'$ with $\varphi' \in \mathcal{O}_{e-\hat{e}_1}^{\text{poly-strong}}$. Then we can apply the same argument as in the case of $\varphi = \langle \alpha \rangle \varphi'$.
 - Case $\psi = \neg \langle \alpha \rangle \varphi'$ with $\varphi' \in \mathcal{O}_{(e \sqcap (e_2, \infty, \infty)) - \hat{e}_3}^{\text{poly-strong}}$. Therefore, $\langle \alpha \rangle \varphi' \in \mathcal{O}_{(e \sqcap (e_2, \infty, \infty)) - \hat{e}_3}^{\text{poly-strong}}$ distinguishes q from p . This time, we can employ the same argument as in the case of $\varphi = \langle \alpha \rangle \varphi'$ to obtain that $(e \sqcap (e_2, \infty, \infty)) - \hat{e}_3 \in \text{Win}_a^{\mathcal{G}_B^\sharp}([q, p])$. The move $[p, q] \xrightarrow{\min_{\{1,2\}, 0, -1} \mathcal{G}_B^\sharp} [q, p]$ and $\text{upd}(\min_{\{1,2\}, 0, -1}, e) = (e \sqcap (e_2, \infty, \infty)) - \hat{e}_3$ justify that $e \in \text{Win}_a^{\mathcal{G}_B^\sharp}([p, q])$. \square

Theorem 4.1. $p \preceq_{\mathcal{O}_e} q$ for $e \in \mathbb{N}^{\text{poly-strong}}$ precisely if the defender wins \mathcal{G}_B^\sharp from $[p, q]$ with budget e .

Thus, we have established that the bisimulation energy game \mathcal{G}_B^\sharp characterizes the equivalences of the polynomial-time strong spectrum $\mathbb{N}^{\text{poly-strong}}$. To

exploit this property to decide equivalences and solve the spectroscopy problem algorithmically, we need one more ingredient: A decision procedure for winning budgets in declining energy games.

4.3 Deciding Energy Games

This section is about how to compute which energy budgets are winning for the attacker (or the defender) at positions of an energy game:

Definition 4.9 (Monotonic Energy Game Winner Problem). In the context of an energy system (\mathcal{E}, \leq) and a finite energy game $\mathcal{G}^\blacktriangleright = (G, G_d, \gamma \rightarrow, upd)$, the *energy game winner problem* goes:

Input Game position $g \in G$.

Output $\text{Win}_a^{\min}(g)$ —the Pareto front of minimal winning budgets for the attacker at g .

We will show how to solve the problem for energy games with monotonic updates. To this end, we will adapt the idea of back-propagating how the defender loses in reachability games from Algorithm 2.1.

The change is that we have to propagate not just attacker wins, but Pareto fronts of attacker-winning energy levels. For this, we have to “invert” the energy update functions. But, as we have seen in Example 4.2, our specific declining energy updates are neither injective nor surjective, and thus cannot be cleanly inverted.

We will see how to tackle these challenges using *Galois connections*. For details, readers are referred to Lemke (2024).

4.3.1 Galois Connections

Galois connections can be thought of as a pair of *monotonic functions that invert each other up to their partial ordering relations*. Another standard intuition is that Galois connections couple a *concrete* and a more *abstract* domain, and this is what the variable naming in the following definition alludes to.

Definition 4.10. A *Galois connection* between two partially ordered sets (C, \leq_C) and (A, \leq_A) is a pair of functions $\alpha: C \rightarrow A$ and $\gamma: A \rightarrow C$ such that, for all $a \in A$ and $c \in C$:

$$\alpha(c) \leq_A a \iff c \leq_C \gamma(a).$$

An illustration of the connection property can be seen in Figure 4.7. (The illustration also illuminates why the α -function is often referred to as “lower adjoint” and γ as “upper adjoint” of a Galois connection.)

Example 4.4. Take the non-negative reals $\mathbb{R}^{\geq 0}$ as concrete domain and the natural numbers \mathbb{N} as abstract domain. Then, as function $\alpha_{\mathbb{R}}: \mathbb{R}^{\geq 0} \rightarrow \mathbb{N}$,




Figure 4.7: Visualization of Galois connections according to Definition 4.10.

take flooring, $x \mapsto \lfloor x \rfloor$, and for the other direction, the identity $\text{id}_{\mathbb{N}}$, $n \mapsto n$. Clearly, $\lfloor x \rfloor \leq n \iff x \leq n$. Therefore, $\alpha_{\mathbb{R}}$ and $\text{id}_{\mathbb{N}}$ form a Galois connection.

Any monotonic function naturally induces a Galois-connected abstraction function.

Lemma 4.3. *The following are equivalent:*

- $\gamma: A \rightarrow C$ is a monotonic function, and $\alpha(c) = \inf\{a \mid c \leq_C \gamma(a)\}$ for all $c \in C$.
- α and γ constitute a Galois connection between C and A .⁶⁷

We will refer to functions derived like $\alpha(c)$ in Lemma 4.3 as “undo” functions. We understand them to generalize inverse functions of monotonic functions as these will always be Galois connected.

Lemma 4.4. *If a function $f: A \rightarrow C$ is surjective and injective, it has a unique inverse function $f^{-1}: C \rightarrow A$ such that $f^{-1} \circ f = \text{id}_A$ and $f \circ f^{-1} = \text{id}_C$.*

If f is monotonic with respect to some partial orders on C and A , then f^{-1} and f must form a Galois connection.

Proof. Because of injectivity and monotonicity, f and f^{-1} , must be (strictly) monotonic. Together with the definition of function inversion, we can reason

$$f(c) \leq_A a \implies f^{-1}(f(c)) \leq_A f^{-1}(a) \implies c \leq_C f^{-1}(a)$$

and

$$c \leq_C f^{-1}(a) \implies f(c) \leq_A f(f^{-1}(a)) \implies f(c) \leq_A a.$$

□

In the following, we will use a \mathcal{O} -symbol to label the adjoint functions that we use to undo energy updates.

Remark 4.3. Definition 3.5 demands monotonicity of our observation languages for a spectrum, $\mathcal{O}_{N \in \mathbb{N}}$. If we construct an abstraction function as in Lemma 4.3 to undo the language selection for a set of formulas Φ , the result would read:

$$\alpha_{\mathcal{O}}(\Phi) := \inf\{N \in \mathbb{N} \mid \Phi \subseteq \mathcal{O}_N\}.$$

The singleton case of this function exactly matches our expressiveness prices of Definition 3.6:

$$\text{expr}(\varphi) := \inf\{N \in \mathbb{N} \mid \varphi \in \mathcal{O}_N\}.$$

4.3.2 The Algorithm

To implement a back-propagation algorithm, we have to assume that we know a function that undoes energy updates, $\text{upd}^{\mathcal{O}}: (\rightarrow) \rightarrow \mathcal{E} \rightarrow \mathcal{E}$ with




Figure 4.8: Illustration of the Galois connection of Example 4.4 between $\mathbb{R}^{\geq 0}$ and \mathbb{N} using $\lfloor \cdot \rfloor$ -function, whose mapping is illustrated by the gray triangles.

⁶⁷ Proposition 4 of Ern  et al. (1993).

$upd^\mathcal{G}(g, g')(e') = \inf\{e \mid e' \leq upd(g, g')(e)\}$. That means, we assume that there is a Galois connection $(upd^\mathcal{G}, upd)$ between energies \mathcal{E} and the domain of updates $\{e \in \mathcal{E} \mid upd(e) \neq \perp\}$.

Algorithm 4.1 can be used to compute the minimal attacker winning budgets.

```

1  def compute_winning_budgets( $\mathcal{G} = (G, G_d, \succ, upd), upd^\mathcal{G}$ ):
2    attacker_win := [ $g \mapsto \emptyset \mid g \in G$ ]
3    todo :=  $\{g \in G_d \mid g \not\mapsto\}$ 
4    while todo  $\neq \emptyset$ :
5      g := some todo
6      todo := todo  $\setminus \{g\}$ 
7      if  $g \in G_a$ :
8        new_attacker_win := Min( $\{upd^\mathcal{G}(g, g')(e') \mid g \mapsto g' \wedge e' \in attacker\_win[g']\}$ )
9      else:
10       new_attacker_win :=  $\{0\}$ 
11       for  $g' \in (g \succ \cdot)$ :
12         new_attacker_win := Min( $\{sup(e_a, upd^\mathcal{G}(g, g')(e')) \mid e_a \in new\_attacker\_win \wedge g \mapsto g' \wedge e' \in attacker\_win[g']\}$ )
13       if new_attacker_win  $\neq$  attacker_win[g]:
14         todo := todo  $\cup (\cdot \succ g)$ 
15     Winamin := attacker_win
16     return Winamin

```

Algorithm 4.1: Algorithm determining the minimal attacker winning budgets Win_a^{\min} of an energy game \mathcal{G} .

The algorithm can be understood as a generalization of Algorithm 2.1, but as positions might need to be revisited more than once, we have to depart from the option counting trick. Whenever we learn about a new minimal energy for the attacker to win, we schedule a position to be (re-)visited. The chain-reaction starts at defender positions with no outgoing moves (lines 3 and 10). By 0 , we denote the minimal element of the energy system \mathcal{E} .


We illustrate the algorithm by running it on the simple energy game of Example 4.1.

Example 4.5. Let us execute Algorithm 4.1 on the \mathbb{N}^2 -game of Example 4.1. The game graph, labeled by the minimal attacker winning energies we find, is reproduced in Figure 4.9.

What we need for the algorithm to run, is some undo function $upd^\mathcal{G}: (\mapsto) \rightarrow \mathcal{E} \rightarrow \mathcal{E}$. We define it as:

$$upd^\mathcal{G}(e') := \sup(e - \vec{u}, 0)$$

We list the steps of the computation in Example 4.5, and the mechanics behind



Figure 4.9: Computed winning budgets Win_a^{\min} on the game in Example 4.5.

them below.

Table 4.1: Steps while solving the game in Example 4.5.

step	g	new_attacker_win	todo
0	-	$g_i \mapsto \emptyset$	g_6
1	g_6	$\{(0, 0)\}$	g_4, g_5
2	g_4	\emptyset	g_5
3	g_5	$\{(0, 1)\}$	g_4
4	g_4	$\{(1, 1)\}$	g_1, g_2, g_3
5	g_1	$\{(3, 2)\}$	g_2, g_3
6	g_2	$\{(0, 2)\}$	g_3, g_1
7	g_3	$\{(1, 1)\}$	g_1
8	g_1	$\{(0, 2), (2, 1)\}$	\emptyset

0. The algorithm begins with g_6 in the todo-set as this is the only position where the defender is stuck.
1. Updating the defender node $g_6 \in G_d$, we set it its minimal attacker-winning budget to 0. This means that the attacker wins here with any budget. As there has been a change, both predecessors g_4 and g_5 are added to the todo. (Their order is undefined; for the sake of the simulation we will always pick the first positions from the todo.)
2. Updating $g_4 \in G_d$ finds no attacker wins so far because one successor $\text{attacker_win}[g_5] = \emptyset$.
3. $g_4 \in G_a$ is updated with the $\text{upd}^\circ((0, -1))((0, 0)) = \sup((0, 0) - (0, -1), \mathbf{0}) = (0, 1)$. This places g_4 in todo again.
4. Now, we can find attacker winning budgets for $g_4 \in G_d$: The two successor options (after undo-update) are $(0, 1)$ and $(1, 0)$, as depicted in Figure 4.10. The unless the attacker at least arrives with energy $\sup((0, 1), (1, 0)) = (1, 1)$, the defender has a way of winning. Effectively, this amounts to intersecting attacker winning budgets of possible successor positions.
5. For $g_1 \in G_a$, it suffices to know a winning budget for one successor. So we determine as a preliminary winning budget $\text{upd}^\circ((-2, -1))((1, 1)) = (3, 2)$.
6. $g_2 \in G_a$ also derives its winning budget from $(1, 1)$. But $(1, 1) - (2, -1) = (-1, 2)$ would be outside the possible energy levels the attacker could carry here and is thus sup-ed into \mathbb{N}^2 , yielding new win $(0, 2)$.
7. $g_3 \in G_a$ just propagates $(1, 1)$ from g_4 .
8. Revisiting $g_1 \in G_a$, we can now combine the three Pareto-fronts of its successors as illustrated in Figure 4.11 as $\text{Min}\{(0, 2), (2, 1), (3, 4)\} = \{(0, 2), (2, 1)\}$. This operation corresponds to taking the union of the

Figure 4.10: Multiplicative combination (dashed red) of successor Pareto fronts (black) at defender position g_4 in step 4.Figure 4.11: Additive combination (dashed red) of successor Pareto fronts (black) at attacker position g_1 in step 8.

(3, 4) of step 5 is thus superseded. As `todo` is empty, the algorithm terminates.

Theorem 4.2. *Assume energies form a well-founded bounded \sqcup -semi-lattice (\mathcal{E}, \leq) . Given a finite-state energy game $\mathcal{G}^\blacktriangleright = (G, G_d, \succrightarrow, upd)$ with monotonic upd^\ominus -functions for all moves and with computable undo function $upd^\ominus(g, g')(e') = \min\{e \mid e' \leq upd(g, g')(e)\}$, Algorithm 4.1 computes the Pareto front of minimal attacker winning budgets, solving the monotonic energy game winner problem (Definition 4.9).*

Proof. Algorithm 4.1 is a work-list variant of a fixed-point iteration computing the least fixed point of attacker winning budgets Win_a according to the inductive characterization of Proposition 4.2.

The “least” here refers to the size of the Win_a -sets characterized by the Win_a^{\min} -Pareto front.

A detailed proof, employing Kleene’s fixed point theorem, can be found in Lemke (2024). Theorem 5 of Lemke (2024) refers to a variant of Algorithm 4.1, where *all* game positions are updated *simultaneously* in each iteration, which aligns more directly with the underlying functor. As is common for graph algorithms, Algorithm 4.1 with updates of single vertices whose neighborhood has changed computes the same result, and usually in a more efficient manner. \square

4.3.3 Complexity

Thanks to Lemke (2024), we can name quite precise bounds for the running time of the algorithm in terms of size of the game and shape of the energy lattice.

Definition 4.11. The *width*, $\text{wdh}_B(b)$, of a (semi-)lattice (B, \leq) below a point b is defined as the maximal size of anti-chains $B' \subseteq \downarrow \{b\}$. (Cf. Definition 3.4.)

The *height*, $\text{hgt}_B(b)$, below b is defined as the size of the maximal size of chains $B' \subseteq \downarrow \{b\}$.

Theorem 4.3. *Consider the following parameters of an energy game problem on $\mathcal{G}^\blacktriangleright = (G, G_d, \succrightarrow, upd)$ where the energies at least form a well-founded bounded \sqcup -semi-lattice (\mathcal{E}, \leq) :*

- t_{sup} , time to calculate the supremum between two elements of \mathcal{E} ,
- t_{\leq} , time to compare two elements of \mathcal{E} ,
- t_{upd^\ominus} , time to compute upd^\ominus on an energy,
- $e_{\mathcal{G}}$, the highest energy that can be achieved by applying permutations of $upd^\ominus(\cdot)$ to $\mathbf{0}$ for $|G| - 1$ times.
- o_{\succrightarrow} , the out-degree of the game graph \succrightarrow .

Then, Algorithm 4.1 computes the minimal attacker winning budgets in

$$O((|G| + \text{hgt}_{\mathcal{E}}(e_{\mathcal{G}})) \cdot |G| \cdot o_{\succrightarrow} \cdot (\text{wdh}_{\mathcal{E}}(e_{\mathcal{G}}))^2 \cdot (t_{\text{upd}^\ominus} + t_{\text{sup}} + \text{wdh}_{\mathcal{E}}(e_{\mathcal{G}}) \cdot t_{\leq})).$$

If $\text{upd}^\mathcal{G}$ is strictly declining, then the following suffices:

$$O(|G|^2 \cdot o_{\rightarrow} \cdot (\text{wdh}_\mathcal{E}(e_\mathcal{G}))^2 \cdot (t_{\text{upd}^\mathcal{G}} + t_{\text{sup}} + \text{wdh}_\mathcal{E}(e_\mathcal{G}) \cdot t_{\leq})).$$

In both cases, the algorithm needs $O(|G| \cdot \text{wdh}_\mathcal{E}(e_\mathcal{G}) \cdot N)$ space to store the Pareto fronts.

Proof. As in Theorem 4.2, we refer to Lemke (2024) for a very detailed proof, where the corresponding fact is proved as Theorem 7.

TODO: Add a proof sketch regardless. \square

4.3.4 Solving Declining Energy Games

Let us instantiate the previous two subsections to declining energy games. All we have to do is instantiate $\text{upd}^\mathcal{G}$ to obtain an algorithm and think about the structure of the energies to tell its complexity.

Definition 4.12. The undo-update function $\text{upd}^\mathcal{G}$ on an N -dimensional declining energy game is defined component-wise for $k \in \{1 \dots N\}$ as:

$$(\text{upd}^\mathcal{G}(u, e'))_k := \max(\{e'_k - u_k \mid u_k \in \{0, -1\}\} \cup \{e'_j \mid \exists D. u_j = \min_D \wedge k \in D\}).$$

TODO: Example

Lemma 4.5. For all updates $u \in \text{Up}$, $\text{upd}^\mathcal{G}(u, \cdot)$ and $\text{upd}(u, \cdot)$ form a Galois connection on the domain $\text{dom}(u)$.

Proof. As $\text{upd}(u, \cdot)$ is monotonic, all that needs to be shown according to Lemma 4.3, is that $\text{upd}^\mathcal{G}(u, e') = \min\{e \mid e' \leq \text{upd}(u, e)\}$.

- Soundness: $\text{upd}^\mathcal{G}(u, e') \in \{e \mid e' \leq \text{upd}(u, e)\}$. This boils down to showing $e'_k \leq (\text{upd}(u, \text{upd}^\mathcal{G}(u, e')))_k$. Consider the cases of u_k .
 - Case $u_k \in \{0, -1\}$. Then, we have to show $e'_k \leq \text{upd}^\mathcal{G}(u, e') + u_k$. Due to $\text{upd}^\mathcal{G}$ being increasing, this must be the case as $e'_k \leq e'_k + u_k$.
 - Case there is D such that $u_k = \min_D$ and $k \in D$. Then, by definition of \min -updates, we have to show $e'_k \leq \min\{(\text{upd}^\mathcal{G}(u, e'))_j \mid j \in D\}$, that is, $e'_k \leq \text{upd}^\mathcal{G}(u, e')_j$ for all $j \in D$. We do so by contradiction: Assume there were a j where this ordering does not hold, meaning $e'_k > \text{upd}^\mathcal{G}(u, e')_j$. As $k \in D$, this is impossible due to the definition of $\text{upd}^\mathcal{G}$ in terms of \max .
- Minimality: For all $e' \leq \text{upd}(u, e)$, we show $\text{upd}^\mathcal{G}(u, e') \leq e$. More specifically, $(\text{upd}^\mathcal{G}(u, e'))_k \leq e_k$. Let us define $d' := \max\{e'_j \mid \exists D. u_j = \min_D \wedge k \in D\}$. Consider the cases of $u_k \in \text{Up}$.

- Case $u_k \in \{0, -1\}$ and $e'_k - u_k \geq d'$. Then $e'_k \leq (\text{upd}(u, e))_k = e_k + u_k$, implying $e'_k - u_k \leq (\text{upd}(u, e))_k = e_k$. By definition, $(\text{upd}^\mathcal{G}(u, e'))_k = \max(e'_k - u_k, d')$. As $e'_k - u_k \geq d'$, we can connect to the previous inequality, getting $(\text{upd}^\mathcal{G}(u, e'))_k = e'_k - u_k \leq e_k$.
- Otherwise, $(\text{upd}^\mathcal{G}(u, e'))_k = d'$. We show $d' \leq e_k$ by contradiction. Assume it were the case that $d' > e_k$. For this to be true there would need to be j and D such that $u_j = \min_D$, $k \in D$, and $e'_j \geq d' > e_k$. By definition of min-update at j , we calculate $(\text{upd}(u, e))_j \leq e_k < e'_j$. But this would contradict the global assumption that $e' \leq \text{upd}(u, e)$. \square

Lemma 4.6. *Algorithm 4.1 solves the energy game problem on an N -dimensional declining energy game $\mathcal{G}^\blacktriangleright = (G, G_d, \mapsto, w)$ in time*

$$O(o_{\mapsto} \cdot |G|^{2 \cdot N} \cdot (N^2 + |G|^{N-1} \cdot N))$$

and space $O(|G|^N \cdot N)$.

Proof. For declining updates, we instantiate the second case of Theorem 4.3:

$$O(|G|^2 \cdot o_{\mapsto} \cdot (\text{wdh}_{\mathcal{E}}(e_{\mathcal{G}}))^2 \cdot (t_{\text{upd}^\mathcal{G}} + t_{\text{sup}} + \text{wdh}_{\mathcal{E}}(e_{\mathcal{G}}) \cdot t_{\leq})).$$

For *declining energy games* as defined in Definition 4.5, we can fill in several blanks:

- t_{sup} , the time to calculate the supremum between two elements of En is in $O(N)$,
- t_{\leq} , time to compare two elements of En is in $O(N)$,
- $t_{\text{upd}^\mathcal{G}}$, time to compute $\text{upd}^\mathcal{G}$ on an energy is in $O(N^2)$,
- $e_{\mathcal{G}}$, the highest energy that can be achieved by applying permutations of $\text{upd}^\mathcal{G}$ to $\mathbf{0}$ for $|G| - 1$ times, is the N -ary vector $(|G| - 1, \dots, |G| - 1)$,
- $\text{wdh}_{\text{En}}(e_{\mathcal{G}})$, the widest anti-chain under $e_{\mathcal{G}}$, can be bounded $O(|G|^{N-1})$.

This yields:

$$\begin{aligned} & O(|G|^2 \cdot o_{\mapsto} \cdot (\text{wdh}_{\text{En}}(e_{\mathcal{G}}))^2 \cdot (t_{\text{upd}^\mathcal{G}} + t_{\text{sup}} + \text{wdh}_{\text{En}}(e_{\mathcal{G}}) \cdot t_{\leq})) \\ &= O(|G|^2 \cdot o_{\mapsto} \cdot (|G|^{N-1})^2 \cdot (N^2 + N + |G|^{N-1} \cdot N)) \\ &= O(o_{\mapsto} \cdot |G|^{2N+2-2} \cdot (N^2 + |G|^{N-1} \cdot N)) \\ &= O(o_{\mapsto} \cdot |G|^{2N} \cdot (N^2 + |G|^{N-1} \cdot N)) \end{aligned}$$

For space complexity, we obtain:

$$\begin{aligned} & O(|G| \cdot \text{wdh}_{\text{En}}(e_{\mathcal{G}}) \cdot N) \\ & O(|G| \cdot |G|^{N-1} \cdot N) \\ & O(|G|^N \cdot N) \end{aligned}$$

\square

Definition 4.13. For N -dimensional energy games, we define the *flattened energies* as $\widehat{\mathbf{En}} := (\{0, 1, \infty\})^N$. A standard energy $e \in \mathbf{En}$ is cast to $\widehat{\mathbf{En}}$ by \widehat{e} where

$$(\widehat{e})_k := \begin{cases} e_k & \text{if } e_k \leq 1 \\ \infty & \text{otherwise} \end{cases}.$$

We denote as *flattened monotonic energy game winner problem* the variant of Definition 4.9 that outputs

$$\widehat{\text{Win}}_a^{\min}(g) := \text{Min}(\text{Win}_a(g) \cap \widehat{\mathbf{En}}).$$

Clearly, Algorithm 4.1 can too be used to solve the flattened version of the problem by adapting upd^{G} to represent all components above 1 by ∞ . Effectively, this decouples the size of Pareto fronts from the game size. In this instance, Lemma 4.6 becomes:

Lemma 4.7. *Algorithm 4.1 solves the flattened energy game problem on an N -dimensional declining energy game $\mathcal{G}^{\blacklozenge} = (G, G_d, \rightsquigarrow, w)$ in time*

$$O(o_{\rightsquigarrow} \cdot |G|^{2N} \cdot (N^2 + 3^{N-1} \cdot N))$$

and in space $O(|G| \cdot 3^{N-1} \cdot N)$.

The bounds we have established are *exponential* with regard to the dimensionality N . But in our use case, we care for energy games of fixed dimensionality. The time bounds are *polynomial* with regard to game graph size. The space bounds behave similarly, but in the *flattened* variant fixed energy makes the space usage even drop to be *linear* in terms of game graph size.

4.3.5 Polynomial Spectroscopy Complexity

We have gathered enough material to have this section prove that the spectroscopy problem for the polynomial-time strong spectrum is indeed solvable in polynomial time, thus justifying the name we chose for the spectrum.

Theorem 4.4. *Given a transition system \mathcal{S} , the spectroscopy problem for the $\mathbf{N}^{\text{poly-strong-spectrum}}$ can be solved in polynomial time with respect to the size of \mathcal{S} .*

Proof. Theorem 4.1 has established that we can solve the spectroscopy problem for the $\mathbf{N}^{\text{poly-strong-spectrum}}$ by deciding the winning budgets of the bisimulation energy game $\mathcal{G}_B^{\blacklozenge}$ on $\mathcal{S} = (\mathcal{P}, \text{Act}, \rightarrow)$. Thus, all that remains to be done is to instantiate the winning budget complexity of Lemma 4.6 for the case $N = 3$ with the size of $\mathcal{G}_B^{\blacklozenge}$ according to Definition 4.7.

The amount of attacker positions is bounded by $|\mathcal{P}|^2$. These positions can initiate up to $|\mathcal{P}| \cdot |\rightarrow|$ simulation challenges, leading to a similarly-bounded amount of defender positions, which again can be left by $|\mathcal{P}| \cdot |\rightarrow|$ moves.

In total, this amounts to $o_{\rightsquigarrow_B}^{\blacklozenge}$ in $O(|\mathcal{P}| \cdot |\rightarrow|)$ and to $|G_B^{\blacklozenge}|$ in $O(|\mathcal{P}|^2)$.

Inserting the parameters in the time bounds of Lemma 4.6 yields:

$$\begin{aligned}
& O(\quad o_{\rightarrow} \quad \cdot \quad |G|^{2N} \quad \cdot \quad (N^2 + |G|^{N-1} \cdot N) \quad) \\
= & O((|\mathcal{P}| \cdot |\rightarrow|) \cdot (|\mathcal{P}|^2)^{2 \cdot 3} \cdot (3^2 + (|\mathcal{P}|^2)^{3-1} \cdot 3) \quad) \\
= & O((|\mathcal{P}| \cdot |\rightarrow|) \cdot |\mathcal{P}|^{12} \cdot |\mathcal{P}|^4 \quad) \\
= & O(|\mathcal{P}|^{17} \cdot |\rightarrow| \quad).
\end{aligned}$$

For space complexity, the approach arrives at $O((|\mathcal{P}| \cdot |\rightarrow|)^3)$. The bound drops to $O((|\mathcal{P}| \cdot |\rightarrow|))$ in the flattened variant of Lemma 4.7. \square

4.4 Discussion

This chapter has shown how to solve the spectroscopy problem for the $\text{N}^{\text{poly-strong}}$ spectrum in polynomial time, thereby deciding all its preorders and equivalences *at once*.

The core ingredient has been to characterize the spectrum's observation languages through an energy game (Note 5, Theorem 4.1) derived from the bisimulation game. The Pareto front view is slightly more general than other characteristic games, which Hüttel & Shukla (1996) for instance use to establish polynomial upper bounds individually for simulation-like notions.

To compute the Pareto fronts, we use a fixed point algorithm (Note 6, Algorithm 4.1). It can be thought of as a generalization of well-known ways to compute shortest distances in graphs:

Remark 4.4. The energy game problem of this chapter can be seen as a *generalization of the shortest path problem* on directed graphs to account for two players as well as for much more general distances and cost functions.

To see how, consider energy games on (\mathbb{Z}, \leq) where only deadlock positions belong to the defender and where all updates constitute one-dimensional integer addition. Then, $e \in \text{Win}_a^{\min}(g)$ iff e is the length of a shortest path from g to some defender position (where negative $(-z)$ -updates stand for what is usually written as a positive edge weight z for graphs). In this instance, Algorithm 4.1 behaves like the well-known Bellman–Ford algorithm for shortest paths.

Algorithm 4.1 contains relevant generalizations and practical improvements compared to the one originally used in Bisping (2023a). Most importantly, the declining energy version loses a dimension of exponentiality. For the polynomial equivalences we arrive at a reasonably polynomial complexity for the spectroscopy instantiation, including reasonable space complexity for a flattened energy lattice.

Knowing that parts of the strong spectrum are PSPACE-hard, the algorithmic complexity indirectly proves that the bisimulation game cannot be used to characterize the whole spectrum. (Unless the polynomial hierarchy miraculously collapses to $\text{PSPACE} = \text{P}$.)

Thus, “we’re gonna need a bigger boat” *game*, at least exponentially-sized, to also cover PSPACE-hard notions.

5 Spectroscopy of the Strong Equivalence Spectrum

i Todo

This section will be filled with contents from [Process Equivalence Problems as Energy Games](#).

5.1 The Strong Spectroscopy Game

5.1.1 The Game

Definition 5.1. For a system $\mathcal{S} = (\mathcal{P}, \text{Act}, \rightarrow)$, the 6-dimensional *spectroscopy energy game* $\mathcal{G}_{\Delta}^{\mathcal{S}}[g_0, e_0] = (G, G_d, \rightsquigarrow, w, g_0, e_0)$ consists of

- *attacker positions* $[p, Q]_a \in G_a$,
- *attacker clause positions* $[p, q]_a \in G_a$,
- *defender conjunction positions* $(p, Q, Q_*)_d \in G_d$,

where $p, q \in \mathcal{P}$ and $Q, Q_* \in 2^{\mathcal{P}}$, and six kinds of moves:

<i>observation</i>	$[p, Q]_a$	$\xrightarrow{-1,0,0,0,0,0}$	$[p', Q']_a$	if $p \xrightarrow{\alpha} p', Q \xrightarrow{\alpha} Q'$,
<i>conjunction</i>	$[p, Q]_a$	$\xrightarrow{0,0,0,0,0,0}$	$(p, Q \setminus Q_*, Q_*)_d$	if $Q_* \subseteq Q$,
<i>conj. revival</i>	$(p, Q, Q_*)_d$	$\xrightarrow{\min\{1,3\}, -1, 0, 0, 0, 0}$	$[p, Q_*]_a$	if $Q_* \neq \emptyset$,
<i>conj. answer</i>	$(p, Q, Q_*)_d$	$\xrightarrow{0, -1, 0, \min\{3,4\}, 0, 0}$	$[p, q]_a$	if $q \in Q$,
<i>positive conjunct</i>	$[p, q]_a$	$\xrightarrow{\min\{1,4\}, 0, 0, 0, 0, 0}$	$[p, \{q\}]_a$,	and
<i>negative conjunct</i>	$[p, q]_a$	$\xrightarrow{\min\{1,5\}, 0, 0, 0, 0, -1}$	$[q, \{p\}]_a$	if $p \neq q$.

A schematic representation can be seen in Figure 5.1.

Lemma 5.1. p_0 and q_0 are bisimilar iff the defender wins \mathcal{G}_{Δ} from $[p_0, \{q_0\}]_a$ with e_0 for every initial energy budget e_0 , that is, if $(\infty, \infty, \infty, \infty, \infty, \infty) \in \text{Win}_d^{\max}([p_0, \{q_0\}]_a)$.⁶⁸

⁶⁸ Proof in Lemma 1 of Bisping (2023b).




Figure 5.1: Schematic spectroscopy game \mathcal{G}_Δ of Definition 5.1.

5.1.2 Correctness

5.1.3 Complexity

5.2 More Specific Games

5.2.1 The Clever Game

5.2.2 Instance Games

5.3 Discussion

- Mention option of non-perfect attacker information as alternative to subset construction
- Difference to Bisping et al. (2022):
 - Lower exponentiality in conjunctions
 - No explicit construction of formulas

6 Spectroscopy for the Weak Silent-Step Spectrum

i Todo

This section will be filled, once the journal version of [One Energy Game for the Spectrum between Branching Bisimilarity and Weak Trace Semantics](#) is ready.

6.1 Lifting to the Spectrum of Stable Equivalences

6.1.1 Stable Behavioral Equivalences

6.1.2 The Adapted Spectroscopy Game

6.2 The Weak Spectrum

6.2.1 Special Weak Equivalences

6.2.2 HML with Silent Behavior

6.2.3 Expressing the Weak Spectrum with Quantities

6.3 The Weak Spectroscopy Game

6.3.1 The Game

6.3.2 Correctness

6.4 Discussion

- Mention to add back in strong observations

7 Implementations

7.1 Prototype Implementation: *equiv.io*

7.2 GPU Implementation: *gpuequiv*

Vogel ([2024](#))

7.3 Other Implementations

7.3.1 Computer Game: *The Spectroscopy Invaders*

Trzeciakiewicz ([2021](#))

7.3.2 CAAL Extension

Ozegowski ([2023](#)) mention Timo

8 Conclusion

There is a rich body of work on *process equivalence checking* and *refinement checking* for programs and specifications. This thesis argues that most of it can be viewed through one unified lens of *finding differences* in behavior.

References

- Adler, R. (2022). *Simulation fehlertoleranter Konsensalgorithmen in Hash.ai* [Bachelor's Thesis]. Technische Universität Berlin.
- Alshukairi, Z. (2022). *Automatisierte Reduktion von reaktiver zu starker Bisimilarität* [Bachelor's Thesis]. Technische Universität Berlin.
- Babai, L. (2016). Graph isomorphism in quasipolynomial time [extended abstract]. *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing*, 684–697. <https://doi.org/10.1145/2897518.2897542>
- Balcázar, J., Gabarro, J., & Santha, M. (1992). Deciding bisimilarity is p-complete. *Formal Aspects of Computing*, 4, 638–648. <https://doi.org/10.1007/BF03180566>
- Bell, C. J. (2013). Certifiably sound parallelizing transformations. In *International conference on certified programs and proofs* (pp. 227–242). Springer.
- Bisping, B. (2018). *Computing coupled similarity* [Master's thesis, Technische Universität Berlin]. https://coupledsim.bbispig.de/bisping_computingCoupledSimilarity_thesis.pdf
- Bisping, B. (2023a). Process equivalence problems as energy games. In C. Enea & A. Lal (Eds.), *Computer aided verification* (pp. 85–106). Springer Nature Switzerland. https://doi.org/10.1007/978-3-031-37706-8_5
- Bisping, B. (2023b). *Process equivalence problems as energy games*. Technische Universität Berlin. <https://doi.org/10.48550/arXiv.2303.08904>
- Bisping, B., Brodmann, P.-D., Jungnickel, T., Rickmann, C., Seidler, H., Stüber, A., Wilhelm-Weidner, A., Peters, K., & Nestmann, U. (2016). Mechanical verification of a constructive proof for FLP. In J. C. Blanchette & S. Merz (Eds.), *Interactive theorem proving* (pp. 107–122). Springer International Publishing. https://doi.org/10.1007/978-3-319-43144-4_7
- Bisping, B., & Jansen, D. N. (2024). One energy game for the spectrum between branching bisimilarity and weak trace semantics. In G. Caltais & C. Di Giusto (Eds.), *Proceedings Combined 31st International Workshop on expressiveness in concurrency and 21st Workshop on structural operational semantics, Calgary, Canada, 9th September 2024* (Vol. 412, pp. 71–88). Open Publishing Association. <https://doi.org/10.4204/EPTCS.412.6>
- Bisping, B., Jansen, D. N., & Nestmann, U. (2022). Deciding All Behavioral Equivalences at Once: A Game for Linear-Time–Branching-Time Spectroscopy. *Logical Methods in Computer Science*, 18(3). [https://doi.org/10.46298/lmcs-18\(3:19\)2022](https://doi.org/10.46298/lmcs-18(3:19)2022)
- Bisping, B., & Montanari, L. (2021). A game characterization for contrasimilarity. In O. Dardha & V. Castiglioni (Eds.), *Proceedings Combined 28th International Workshop on expressiveness in concurrency and 18th Workshop*

- on structural operational semantics (Vol. 339, pp. 27–42). Open Publishing Association. <https://doi.org/10.4204/EPTCS.339.5>
- Bisping, B., & Montanari, L. (2023). Coupled similarity and contrasimilarity, and how to compute them. *Arch. Formal Proofs*, 2023. https://www.isa-afp.org/entries/CoupledSim_Contrasim.html
- Bisping, B., & Montanari, L. (2024). Characterizing contrasimilarity through games, modal logic, and complexity. *Inf. Comput.*, 300, 105191. <https://doi.org/10.1016/J.IC.2024.105191>
- Bisping, B., & Nestmann, U. (2019). Computing coupled similarity. In T. Vojnar & L. Zhang (Eds.), *Tools and algorithms for the construction and analysis of systems: TACAS* (Vol. 11427, pp. 244–261). Springer. https://doi.org/10.1007/978-3-030-17462-0_14
- Bisping, B., & Nestmann, U. (2021). A game for linear-time–branching-time spectroscopy. In J. F. Groote & K. G. Larsen (Eds.), *Tools and algorithms for the construction and analysis of systems: TACAS* (Vol. 12651, pp. 3–19). Springer. https://doi.org/10.1007/978-3-030-72016-2_1
- Bisping, B., Nestmann, U., & Peters, K. (2020). Coupled similarity: The first 32 years. *Acta Informatica*, 57(3–5), 439–463. <https://doi.org/10.1007/s00236-019-00356-4>
- Bulik, J. (2021). *Statically analysing inter-process communication in Elixir programs* [Bachelor's Thesis, Technische Universität Berlin]. <https://gitlab.com/MrGreenTea/stille-post>
- Duong, T. M. (2022). *Developing an educational tool for linear-time–branching-time spectroscopy* [Bachelor's Thesis]. Technische Universität Berlin.
- England, J. (2021). *HML synthesis of distinguished processes* [Bachelor's Thesis]. Technische Universität Berlin.
- Erné, M., Koslowski, J., Melton, A., & Strecker, G. E. (1993). A primer on galois connections. *Annals of the New York Academy of Sciences*, 704(1), 103–125. <https://doi.org/10.1111/j.1749-6632.1993.tb52513.x>
- Ésik, Z., Fahrenberg, U., Legay, A., & Quaas, K. (2013). Kleene algebras and semimodules for energy problems. In D. Van Hung & M. Ogawa (Eds.), *Automated technology for verification and analysis* (pp. 102–117). Springer International Publishing. https://doi.org/10.1007/978-3-319-02444-8_9
- Fahrenberg, U., Juhl, L., Larsen, K. G., & Srba, J. (2011). Energy games in multiweighted automata. In A. Cerone & P. Pihlajasaari (Eds.), *Theoretical aspects of computing – ICTAC 2011* (pp. 95–115). Springer. https://doi.org/10.1007/978-3-642-23283-1_9
- Gale, D., & Stewart, F. M. (1953). Infinite games with perfect information. In H. W. Kuhn & A. W. Tucker (Eds.), *Contributions to the theory of games, volume II* (pp. 245–266). Princeton University Press. <https://doi.org/10.1515/9781400881970-014>
- Glabbeek, R. J. van. (1990). The linear time–branching time spectrum: Extended abstract. In J. C. M. Baeten & J. W. Klop (Eds.), *CONCUR'90* (Vol. 458, pp. 278–297). Springer. <https://doi.org/10.1007/BFb0039066>
- Glabbeek, R. J. van. (1993). The linear time–branching time spectrum II: The

- semantics of sequential systems with silent moves; extended abstract. In E. Best (Ed.), *CONCUR'93* (Vol. 715, pp. 66–81). Springer. https://doi.org/10.1007/3-540-57208-2_6
- Glabbeek, R. J. van. (2001). The linear time–branching time spectrum I: The semantics of concrete, sequential processes. In J. A. Bergstra, A. Ponse, & S. A. Smolka (Eds.), *Handbook of process algebra* (pp. 3–99). Elsevier. <https://doi.org/10.1016/B978-044482830-9/50019-9>
- Göthel, T. (2012). *Mechanical verification of parameterized real-time systems* [PhD thesis, Berlin Institute of Technology]. <https://doi.org/10.14279/depositonce-3248>
- Groote, J. F., & Martens, J. (2024). *A quadratic lower bound for simulation*. <https://doi.org/10.48550/arXiv.2411.14067>
- Groote, J. F., Martens, J., & Vink, Erik. P. de. (2023). Lowerbounds for bisimulation by partition refinement. *Logical Methods in Computer Science, Volume 19, Issue 2*. [https://doi.org/10.46298/lmcs-19\(2:10\)2023](https://doi.org/10.46298/lmcs-19(2:10)2023)
- Hauschild, J. (2023). *Nonlinear counterfactuals in Isabelle/HOL* [Bachelor's Thesis, Technische Universität Berlin]. <https://github.com/johanneshauschild/NonlinearCounterfactuals>
- Hennessy, M., & Milner, R. (1980). On observing nondeterminism and concurrency. In J. W. de Bakker & J. van Leeuwen (Eds.), *Automata, languages and programming* (Vol. 85, pp. 299–309). Springer. https://doi.org/10.1007/3-540-10003-2_79
- Hüttel, H., & Shukla, S. (1996). *On the complexity of deciding behavioural equivalences and preorders. A survey* (BRICS RS-96-39H). University of Aarhus. <https://doi.org/10.7146/brics.v3i39.20021>
- Kruskal, J. B. (1972). The theory of well-quasi-ordering: A frequently discovered concept. *Journal of Combinatorial Theory, Series A*, 13(3), 297–305. [https://doi.org/10.1016/0097-3165\(72\)90063-5](https://doi.org/10.1016/0097-3165(72)90063-5)
- Kučera, A., & Esparza, J. (1999). A logical viewpoint on process-algebraic quotients. In J. Flum & M. Rodríguez-Artalejo (Eds.), *Computer science logic: CSL* (Vol. 1683, pp. 499–514). Springer. https://doi.org/10.1007/3-540-48168-0_35
- Kupferman, O., & Shamash Halevy, N. (2022). Energy Games with Resource-Bounded Environments. In B. Klin, S. Lasota, & A. Muscholl (Eds.), *33rd international conference on concurrency theory (CONCUR 2022)* (Vol. 243, pp. 19:1–19:23). Schloss Dagstuhl – Leibniz-Zentrum für Informatik. <https://doi.org/10.4230/LIPIcs.CONCUR.2022.19>
- Kurzan, T. (2024). *Implementierung eines Contrasimilarity-Checkers für mCRL2* [Bachelor's Thesis]. Technische Universität Berlin.
- Lê, H. N. (2020). *Implementing coupled similarity as an automated checker for mCRL2* [Bachelor's Thesis]. Technische Universität Berlin.
- Lemke, C. (2024). *A formal proof of decidability of multi-weighted declining energy games* [Master's thesis]. Technische Universität Berlin.
- Lönne, B. (2023). *An educational computer game about counterfactual truth conditions* [Bachelor's Thesis, Technische Universität Berlin].

- <https://github.com/Brunololos/counterfactuals-demo>
- Mattes, K. P. P. (2024). *Measuring expressive power of HML formulas in Isabelle/HOL* [Bachelor's Thesis, Technische Universität Berlin]. https://github.com/ekeln/BA_formula_prices/raw/main/output/outline.pdf
- Milner, R. (1989). *Communication and concurrency*. Prentice-Hall, Inc.
- Montanari, L. (2021). *Kontrasimulation als Spiel* [Bachelor's Thesis, Technische Universität Berlin]. <https://github.com/luisamontanari/ContrasimGame>
- Nestmann, U., & Pierce, B. C. (2000). Decoding choice encodings. *Information and Computation*, 163(1), 1–59. <https://doi.org/10.1006/inco.2000.2868>
- Ozegowski, F. (2023). *Integration eines generischen äquivalenzprüfers in CAAL* [Bachelor's Thesis, Technische Universität Berlin]. <https://github.com/Fabian-O01/CAAL>
- Paige, R., & Tarjan, R. E. (1987). Three partition refinement algorithms. *SIAM Journal on Computing*, 16(6), 973–989.
- Peacock, D. A. (2020). *Process equivalences as a video game* [Bachelor's Thesis, Technische Universität Berlin]. https://drive.google.com/drive/folders/19YuOCXI_7CEB4p9fgr6kDKCTbFwx2VU9
- Pohlmann, M. (2021). *Reducing strong reactive bisimilarity to strong bisimilarity* [Bachelor's Thesis, Technische Universität Berlin]. <https://maxpohlmann.github.io/Reducing-Reactive-to-Strong-Bisimilarity/thesis.pdf>
- Ranzato, F., & Tapparo, F. (2010). An efficient simulation algorithm based on abstract interpretation. *Information and Computation*, 208(1), 1–22. <https://doi.org/10.1016/j.ic.2009.06.002>
- Reichert, J. M. (2020). *Visualising and model checking counterfactuals* [Bachelor's Thesis]. Technische Universität Berlin.
- Roscoe, A. W. (2009). Revivals, stuckness and the hierarchy of CSP models. *Journal of Logic and Algebraic Programming*, 78(3), 163–190. <https://doi.org/10.1016/j.jlap.2008.10.002>
- Sandt, E. (2022). *A video game about reactive bisimilarity* [Bachelor's Thesis, Technische Universität Berlin]. <https://github.com/eloinoel/ReactiveBisimilarityGame>
- Sangiorgi, D. (2012). *Introduction to bisimulation and coinduction*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511777110>
- Stirling, C. (1996). Modal and temporal logics for processes. In F. Moller & G. Birtwistle (Eds.), *Logics for concurrency: Structure versus automata* (Vol. 1043, pp. 149–237). Springer. https://doi.org/10.1007/3-540-60915-6_5
- Stöcker, V. (2024). *Higher-order diadic μ -calculus—an efficient framework for checking process equivalences?* [Bachelor's Thesis]. Technische Universität Berlin.
- Trzeciakiewicz, M. (2021). *Linear-time–branching-time spectroscopy as an educational web browser game* [Bachelor's Thesis, Technische Universität Berlin]. <https://github.com/Marii19/the-spectroscopy-invaders>
- Vogel, G. (2024). *Accelerating process equivalence energy games using We-*

- bGPU* [Bachelor's Thesis, Technische Universität Berlin]. <https://github.com/Gobbel2000/gpuequiv>
- Voorhoeve, M., & Mauw, S. (2001). Impossible futures and determinism. *Information Processing Letters*, 80(1), 51–58. [https://doi.org/10.1016/S0020-0190\(01\)00217-4](https://doi.org/10.1016/S0020-0190(01)00217-4)
- Wittig, T. (2020). *Charting the jungle of process calculi encodings* [Bachelor's Thesis, Technische Universität Berlin]. <https://concurrency-theory.org/process-jungle/>
- Wortmann, J. K., Olesen, S. R., & Enevoldsen, S. (2015). *CAAL 2.0 – recursive HML, distinguishing formulae, equivalence collapses and parallel fixed-point computations* (Vol. DES103F15). Aalborg University. https://caal.cs.aau.dk/docs/CAAL2_RDEP.pdf
- Wrusch, M. (2020). *Ein Computerspiel zum Erlernen von Verhaltensäquivalenzen* [Bachelor's Thesis]. Technische Universität Berlin.
- Zielonka, W. (1998). Infinite games on finitely coloured graphs with applications to automata on infinite trees. *Theoretical Computer Science*, 200(1), 135–183. [https://doi.org/https://doi.org/10.1016/S0304-3975\(98\)00009-7](https://doi.org/https://doi.org/10.1016/S0304-3975(98)00009-7)