8 Soundness and completeness

In this chapter we will see how to find sound and complete derivation systems for the formulas in ML_Λ that are *valid* in the class of all coalgebras (of the appropriate type). In a slogan, we will show that

completeness of coalgebraic logic is determined at the one-step level,

in the sense that we may transform any one-step sound and complete derivation system $\mathbf{D}$, just by adding some appropriate propositional rules, into a system $\mathbf{D}^+$ that is sound and complete for the validities in the full language ML_Λ. We will prove this result both in the setting of sequent calculi and that of Hilbert-style derivation systems.

Throughout this chapter we fix a set functor T, a modal signature Λ for T, and a countably infinite set A of variables. For convenience we will assume that T is standard.

Remark 8.1 We also assume that T is not the trivial functor $K_\varnothing$ which maps every set to the empty set, and every function to the empty map (which is the identity arrow on the empty set). It can be shown that $K_\varnothing$ admits only the empty coalgebra, and for this reason any coalgebraic modal logic for this functor will be trivial.

8.1 Sequent calculi

In this section we will see, in the framework of sequent calculi, how we can transfer soundness and completeness of a derivation system from the one-step level to the level of the full coalgebraic modal language.

8.1.1 Sequent systems

To recall the setting, given a one-step derivation system $\mathbf{G}$, we obtain a system $\mathbf{G}^+$ by adding some (fixed) propositional rules for classical logic. For simplicity we restrict the propositional connectives to the (expressively complete) set $\{\top, \neg, \wedge\}$.

Definition 8.2 Consider the following propositional (axioms and) rules:

$$\frac{}{a, \neg a}\,\mathsf{Ax0} \qquad \frac{}{\top}\,\mathsf{Ax1} \qquad \frac{a}{\neg\neg a}\,\mathsf{R}_{\neg\neg} \qquad \frac{a \quad b}{a \wedge b}\,\mathsf{R}_{\wedge} \qquad \frac{\neg a, \neg b}{\neg(a \wedge b)}\,\mathsf{R}_{\neg\wedge}$$

Given a one-step derivation system $\mathbf{G}$, we define $\mathbf{G}^+$ as its extension with the above propositional rules. ◁

We define ML_Λ-*instances* of a derivation rule in the obvious way. In particular, instances of the above propositional rules look as follows:

$$\frac{}{\Gamma, \varphi, \neg\varphi}\,\mathsf{Ax0} \qquad \frac{}{\Gamma, \top}\,\mathsf{Ax1} \qquad \frac{\Gamma, \varphi}{\Gamma, \neg\neg\varphi}\,\mathsf{R}_{\neg\neg} \qquad \frac{\Gamma, \varphi \quad \Gamma, \psi}{\Gamma, \varphi \wedge \psi}\,\mathsf{R}_{\wedge} \qquad \frac{\Gamma, \neg\varphi, \neg\psi}{\Gamma, \neg(\varphi \wedge \psi)}\,\mathsf{R}_{\neg\wedge}$$

with Γ an arbitrary set of ML_Λ-formulas. In the case of a one-step rule $\mathsf{d} = \Delta_1, \ldots, \Delta_{n-1}/\Delta_0$, any instance of d can be represented as a rule

$$\frac{\Delta_1[\tau] \quad \cdots \quad \Delta_{n-1}[\tau]}{\Gamma, \Delta_0[\tau]}\,\mathsf{d}$$

where $\tau : A \to \mathrm{ML}_\Lambda$ is a substitution, and $\Gamma \subseteq \mathrm{ML}_\Lambda$ is a set of context formulas. As in the previous chapter, the difference between instances of propositional and one-step rules is that in the case of a one-step rule the context only appears in the conclusion, not in the premises.

Formally we will define rule instances as follows.

Definition 8.3 An *instance* of a derivation rule $\mathsf{d} \in \mathbf{G}^+$ is a triple $(\Theta, \mathsf{d}, \tau)$ such that $\mathsf{conc}(d)[\tau] \subseteq \Theta$. The *context* of such a rule instance is the set $\Theta \setminus \mathsf{conc}(d)[\tau]$. ◁

Example 8.4 We look at standard modal logic (that is, $T = P$), taking the diamond operator $\Diamond$ as the single modality. Let $\mathbf{M}_\Diamond$ be the derivation system consisting of all modal derivation rules of the form

$$\frac{\neg a_0, a_1, \ldots, a_n}{\neg\Diamond a_0, \Diamond a_1, \ldots, \Diamond a_n}\,\mathsf{m}_n$$

In Example 7.38 we saw that this system is one-step complete. It will follow from the results in this section that the system $\mathbf{M}_\Diamond^+$ is a sound and complete derivation system for the set of all valid modal formulas in this language.

The notion of a $\mathbf{G}^+$-*derivation* and its *assumptions* are defined in the obvious way. A $\mathbf{G}^+$-*proof* is a finite $\mathbf{G}^+$-derivation without assumptions; that is, all leaves are labelled with axioms. For completeness we provide the details below.

Definition 8.5 Let $\mathbf{G}$ be a one-step derivation system. A *derivation* in $\mathbf{G}^+$ is a structure (W, C, L, r), where (W, C, r) is a well-founded tree with root r, and L is a map *labelling* every node w of W with an ML_Λ-sequent $L(w)$, in such a way that, for every inner node $t \in W$, the pair $\big(\{L(u) \mid u \in C(t)\}, L(t)\big)$ is an instance of a derivation rule in $\mathbf{G}^+$.

Given an $\mathbf{G}^+$-derivation $\mathcal{D} = (W, C, L, r)$, we call a sequent Σ an *assumption* of $\mathcal{D}$ if $\Sigma = L(t)$ for some *leaf* t, but Σ is not an instance of an axiom; we denote the set of assumptions of $\mathcal{D}$ by $\mathsf{Ass}(\mathcal{D})$. We refer to the sequent $L(r)$ labelling the root of the tree as the *result* of $\mathcal{D}$.

A sequent $\Sigma \in 1\mathrm{ML}_\Lambda(A)$ is *derivable* from a collection $\mathcal{A}$ of sequents in a derivation system $\mathbf{G}^+$, notation: $\mathcal{A} \vdash_{\mathbf{G}^+} \Sigma$ if Σ is the result of an $\mathbf{G}^+$-derivation $\mathcal{D}$ with $\mathsf{Ass}(\mathcal{D}) \subseteq \mathcal{A}$. If $\varnothing \vdash_{\mathbf{G}^+} \Sigma$ we simply write $\vdash_{\mathbf{G}^+} \Sigma$ and we say that Σ is *derivable*. $\lhd$

Example 8.6 Here is an example of an $\mathbf{M}_\Diamond^+$-derivation:

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{}{p,\neg p,\neg\neg q}\,\mathsf{Ax0} \qquad \dfrac{}{\neg q,\neg p,\neg\neg q}\,\mathsf{Ax0}}{p\wedge\neg q,\neg p,\neg\neg q}\,\mathsf{R}_\wedge}{\Diamond(p\wedge\neg q),\Diamond\neg p,\neg\Diamond\neg q,r}\,\mathsf{m}_2 \qquad \dfrac{\dfrac{\dfrac{\dfrac{}{\top}\,\mathsf{Ax1}}{\neg\neg\top}\,\mathsf{R}_{\neg\neg}}{\neg\Diamond\neg\top}\,\mathsf{m}_0}{\Diamond(p\wedge\neg q),\Diamond\neg p,\neg\Diamond\Diamond\neg\top,r}\,\mathsf{m}_0}{\Diamond(p\wedge\neg q),\Diamond\neg p,\neg\Diamond\neg q\wedge\neg\Diamond\Diamond\neg\top,r}\,\mathsf{R}_{\neg\wedge}}{\neg\neg\Diamond(p\wedge\neg q),\Diamond\neg p,\neg\Diamond\neg q\wedge\neg\Diamond\Diamond\neg\top,r}\,\mathsf{R}_{\neg\neg}}{\neg\Diamond\neg\Diamond(p\wedge\neg q),\Diamond\Diamond\neg p,\Diamond(\neg\Diamond\neg q\wedge\neg\Diamond\Diamond\neg\top),r,\Diamond r}\,\mathsf{m}_3}{\neg\Diamond\neg\Diamond(p\wedge\neg q),\neg\neg\Diamond\Diamond\neg p,\Diamond(\neg\Diamond\neg q\wedge\neg\Diamond\Diamond\neg\top),r,\Diamond r}\,\mathsf{R}_{\neg\neg}}{\neg(\Diamond\neg\Diamond(p\wedge\neg q)\wedge\neg\Diamond\Diamond\neg p),\Diamond(\neg\Diamond\neg q\wedge\neg\Diamond\Diamond\neg\top),r,\Diamond r}\,\mathsf{R}_{\neg\wedge}$$

The main result of this section is the following. Recall that we read sequents *disjunctively*; hence, a sequent Σ is *valid* if the formula $\bigvee\Sigma$ is valid (that is, true at every state in every model).

Theorem 8.7 (Soundness and Completeness) *Let* $\mathbf{G}$ *be a set of modal derivation rules which is one-step sound and complete. Then the system* $\mathbf{G}^+$ *is sound and complete for the set of valid* ML_Λ*-sequents. That is, for any sequent* $\Sigma \in \mathsf{Seq}(\mathrm{ML}_\Lambda)$ *the following equivalence holds:*

$$\Sigma \textit{ is valid iff it is derivable in } \mathbf{G}^+. \tag{68}$$

8.1.2 The derivability game

We will prove Theorem 8.7 by *game-theoretic* means. That is, with $\mathbf{G}$ we will associate a *derivability game* $\mathcal{D}_\mathbf{G}@\Sigma$, with two players: *Prover* (P, female) and *Builder* (B, male). The aim of Prover in this game is to find a derivation for Σ, while Builder tries to refute Σ by constructing a *countermodel*, that is, a pointed model where all formulas in Σ are false. The derivability game is played on a board that consists of two kinds of positions:

- Basic positions are given as ML_Λ-sequents, and they are *owned* by Prover. At such a position she has to pick an instance of a rule in $\mathbf{G}^+$ that is applicable to the sequent.
- Secondary positions are given as ML_Λ-instances of derivation rules in $\mathbf{G}^+$ and these positions are *owned* by Builder. At a position of this kind, he has to choose a premise of the rule instance, and in that way he picks the next basic position of the match.

A match of the game $\mathcal{D}_{\mathbf{G}}@\Sigma$ thus proceeds by Prover and Builder successively moving a token from one position to the next, starting from the position Σ. Such a match continues until one of the players gets stuck.

Formally, the game consists of a collection of positions, that provide the vertices of a directed graph (Pos, E). The idea of the game is that, starting from the initial position Σ, the players move a token around the graph, from one position to another, following the edge relation E. Each position is *owned* by one of the players. This ownership relation indicates that, when the token is situated at a position, it is the task of its owner to move the token to a neighbouring position. Here are the formal details of the nature of positions and the edge relation linking them.

- Positions of the form Θ, for some ML_Λ-sequent Θ, are called basic and belong to Prover. She needs to pick a rule $\mathsf{d} \in \mathbf{G}^+$ and a substitution $\tau : A \to \mathrm{ML}_\Lambda$ such that $\mathsf{conc}(\mathsf{d})[\tau] \subseteq \Theta$. Doing so she moves the token to the position $(\Theta, \mathsf{d}, \tau)$.
- Positions of the form $(\Theta, \mathsf{d}, \tau)$, with $\mathsf{conc}(\mathsf{d})[\tau] \subseteq \Theta$, belong to Builder. We think of Θ as the conclusion of the rule instance induced by d and τ and we partition the sequent as $\Theta = \mathsf{conc}(\mathsf{d})[\tau] \uplus \Gamma$, where $\mathsf{conc}(\mathsf{d})[\tau]$ and Γ consist of, respectively, the *principal* and *context* formulas in Θ. It is now Builder's task to choose a premise Δ of d; in doing so he picks the next position Θ' in the match as follows:

$$\Theta' := \begin{cases} \Delta[\tau] \cup \Gamma & \text{if } \mathsf{d} \text{ is a propositional rule,} \\ \Delta[\tau] & \text{if } \mathsf{d} \text{ is a modal rule.} \end{cases}$$

If the token arrives at a position without neighbours (that is, a node $p \in \mathsf{Pos}$ such that $E(v) = \varnothing$), we say that the position's owner *gets stuck*, and when this happens the player immediately looses the match of the game. Both players might find themselves in this situation:

- Prover gets stuck at a sequent position if there is no rule applicable to the sequent. Examples of such sequents include $\{p, \neg q\}$, $\{\neg\top\}$, and (in the case of the derivation system $\mathbf{M}^+_\Diamond$ of Example 8.4) $\{p, \Diamond p, \Diamond\neg p, \Diamond\neg q\}$ or $\{\Diamond\top\}$.
- Builder gets stuck at a position of the form $(\Theta, \mathsf{d}, \tau)$ if d is a rule without premises, that is, if Θ is an instance of an axiom.

Position $p \in \mathsf{Pos}$	Owner $O(p)$	Admissible moves $E(p) \subseteq \mathsf{Pos}$
Θ	P	$\{(\Theta, \mathsf{d}, \tau) \mid \mathsf{d} \in \mathbf{G}^+, \tau \in \mathrm{ML}_\Lambda^A, \mathsf{conc}(\mathsf{d})[\tau] \subseteq \Theta\}$
$(\Theta, \mathsf{d}, \tau)$, d propositional	B	$\{\Delta[\tau] \cup (\Theta \setminus \mathsf{conc}(\mathsf{d})[\tau]) \mid \Delta \in \mathsf{Prem}(\mathsf{d})\}$
$(\Theta, \mathsf{d}, \tau)$, d one-step	B	$\{\Delta[\tau] \mid \Delta \in \mathsf{Prem}(\mathsf{d})\}$

Table 2: Derivability game for a derivation system $\mathbf{G}^+$

Definition 8.8 The *derivability game* $\mathcal{D}_\mathbf{G}$ is defined as the two-player game (Pos, O, E) of which the game graph (Pos, E) and owner function $O : \mathsf{Pos} \to \{P, B\}$ are given in Table 2. For each player $Q \in \{P, B\}$, we denote by $\mathsf{Pos}_Q := O^{-1}(\mathsf{Pos})$ the set of positions they own. Given an ML_Λ-sequent Σ, we let $\mathcal{D}_\mathbf{G}@\Sigma$ denote the version of the game $\mathcal{D}_\mathbf{G}$ that is initialised at Σ. ◁

Before we show how this derivability game can be used to proved soundness and completeness of the derivation system $\mathbf{G}^+$, we prove two fundamental results about the game itself.

First we show that $\mathcal{D}_\mathbf{G}$-matches always end after finitely many moves; that is, all matches are finite. A *match* of this game is a sequence of positions that arise from Prover and Builder successively moving a token from one position to another, following the rules of the game. Formally then, a match is nothing but a path through the game graph (Pos, E). Such a match π is *full* if it is either infinite or it is finite and its final position $\mathsf{last}(\pi)$ has no successors in the game graph. In the latter case, the owner of this final position got stuck and so their opponent is declared to be the winner of π. Matches that are not full are called *partial*. Matches of the initialised game $\mathcal{D}_\mathbf{G}@\Sigma$ must start at Σ.

To establish that all $\mathcal{D}_\mathbf{G}$-matches are in fact finite, we need some complexity measures for formulas and sequents.

Definition 8.9 The *modal depth* of ML_Λ-formulas is given by the following formula induction:

$$
\begin{array}{lcll}
\mathtt{md}(\varphi) & := & 0 & \text{if } \varphi \text{ is atomic} \\
\mathtt{md}(\neg\varphi) & := & \mathtt{md}(\varphi) & \\
\mathtt{md}(\varphi_0 \wedge \varphi_1) & := & \max(\mathtt{md}(\varphi_0), \mathtt{md}(\varphi_1)) & \\
\mathtt{md}(\heartsuit_\lambda(\varphi_1, \ldots, \varphi_n)) & := & 1 + \max\{\mathtt{md}(\varphi_1), \ldots, \mathtt{md}(\varphi_n)\} &
\end{array}
$$

The modal depth of a sequent Φ is defined as $\mathtt{md}(\Phi) := \max\{\mathtt{md}(\varphi) \mid \varphi \in \Phi\}$.

We define the *propositional complexity* of ML_Λ-formulas by the following induction:

$$
\begin{array}{lcll}
\mathtt{pc}(\varphi) & := & 0 & \text{if } \varphi \text{ is atomic} \\
\mathtt{pc}(\neg\varphi) & := & 1 + \mathtt{pc}(\varphi) & \\
\mathtt{pc}(\varphi_0 \wedge \varphi_1) & := & 1 + \mathtt{pc}(\varphi_0) + \mathtt{pc}(\varphi_1) & \\
\mathtt{pc}(\heartsuit_\lambda(\varphi_1, \ldots, \varphi_n)) & := & 0 &
\end{array}
$$

The propositional complexity of a sequent Φ is given as $\mathtt{pc}(\Phi) := \sum_{\varphi \in \Phi} \mathtt{pc}(\varphi)$. ◁

Proposition 8.10 *$\mathcal{D}_\mathbf{G}$ is a finite game; that is, all matches of $\mathcal{D}_\mathbf{G}$ are finite.*

Proof. We leave it as an exercise for the reader to prove this result, using the complexity
T measures of Definition 8.9. QED

As our second result on the derivability game we will show that it enjoys the property of *positional determinacy.* To explain this notion we need to discuss one of the key concepts in game theory, viz., that of a player having a *winning strategy* in an (initialised) game. Intuitively, a winning strategy for player $Q \in \{P, B\}$ in $\mathcal{D}_{\mathbf{G}}@\Sigma$ is a way of playing for Q which guarantees, no matter the response of Q's opponent $\overline{Q}$, that Q will win any match starting from Σ, as long as they stick to the given strategy. The winning region for Q in a game consists of the set of those positions for which Q has a winning strategy.

▶ `say something about formalising these concepts in general`

In the case of finite games, however, we may always assume that winning strategies are *positional.*

Definition 8.11 A *positional strategy* for player $Q \in \{P, B\}$ in $\mathcal{D}_{\mathbf{G}}$ is a partial map $f : \mathsf{Pos}_Q \overset{\circ}{\to} \mathsf{Pos}$ such that $f(p) \in E(p)$, for every position p in the domain of f. A match $\pi = p_0 \cdots p_n$ is *guided* by a positional strategy f for Q if $p_{i+1} = f(p_i)$ for every $i < n$ such that $p_i \in \mathsf{Pos}_Q$.

A positional strategy is *winning for Q* at position $p \in \mathsf{Pos}$ if for every f-guided match π such that $\mathsf{first}(\pi) = p$ and $\mathsf{last}(\pi) \in \mathsf{Pos}_Q$ we have $\mathsf{last}(\pi) \in \mathsf{Dom}(f)$. A positional strategy is *winning* for a region $P \subseteq \mathsf{Pos}$ if $P \subseteq \mathsf{Dom}(f)$ and f is winning for every position $p \in P$. ◁

Observe that a positional strategy f being winning for a player at first glance only ensures that the player will never get stuck when playing f; but since in this game loosing means getting stuck, simply surviving actually means being guaranteed to win.

The main property of the derivability game that we shall need below is the following.

Theorem 8.12 (Positional Determinacy) *The set* Pos *of positions in the game* $\mathcal{D}_{\mathbf{G}}$ *can be partioned as*

$$\mathsf{Pos} = W_P \uplus W_B,$$

in such a way that each player Q has a positional strategy which is winning for the region W_Q.

▶ `Proof (sketch) to be supplied`

It is not hard to see that the partition mentioned in Theorem 8.12 is in fact unique. This inspries and justifies the following definition.

Definition 8.13 We let Win_P and Win_B denote the sets given in Theorem 8.12. For a given player $Q \in \{P, B\}$, we refer to Win_Q as the *winning region of Q.* ◁

Note that the winning region for player Q may contain both positions that are owned by Q and positions that are owned by Q's opponent.

8.1.3 Soundness and completeness

With the definition and basic properties of the derivability game in place, we will now show how $\mathcal{D}_{\mathbf{G}}$ can be used to prove the soundness and completeness of the derivation system $\mathbf{G}^+$. The intuitions underlying the proof are quite straightforward: a winning strategy for Prover in the initialised game $\mathcal{D}_{\mathbf{G}}@\Sigma$ corresponds to a *proof* of Σ in $\mathbf{G}^+$, while a winning strategy for Builder can be seen as a *countermodel* for Σ, or at least as a blueprint for constructing such a model. Soundness and completeness of $\mathbf{G}^+$ then follow from the determinacy of the derivability game: since either Prover or Builder must have a winning strategy in $\mathcal{D}_{\mathbf{G}}@\Sigma$, we are guaranteed to find either a proof or a countermodel for Σ. Here is a more formal proof.

Proof of Theorem 8.7. Let $\mathbf{G}$ be some one-step sequent system which is one-step sound and complete, and let $\Sigma \in \mathsf{Seq}(\mathrm{ML}_\Lambda)$ be some sequent. Consider the following chain of equivalent statements:

$$\begin{array}{rlr} \Sigma \text{ is derivable in } \mathbf{G}^+ & \text{iff } \Sigma \in \mathsf{Win}_P(\mathcal{D}_{\mathbf{G}}) & \text{(Proposition 8.14)} \\ & \text{iff } \Sigma \notin \mathsf{Win}_B(\mathcal{D}_{\mathbf{G}}) & \text{(Theorem 8.12)} \\ & \text{iff } \Sigma \text{ does not have a countermodel} & \text{(Proposition 8.15)} \\ & \text{iff } \Sigma \text{ is valid} & \text{(obvious)} \end{array}$$

This obviously proves the Theorem. QED

Thus, all that is left is to prove the two Propositions below. The proof of the first Proposition is left to the reader.

Proposition 8.14 *Let Σ be some ML_Λ-sequent. Then Prover has a winning strategy in $\mathcal{D}_{\mathbf{G}}@\Sigma$ iff Σ has a $\mathbf{G}^+$-proof.* T

The heart of the completeness proof is the Proposition below.

Proposition 8.15 *Let Σ be some ML_Λ-sequent. Then Builder has a winning strategy in $\mathcal{D}_{\mathbf{G}}@\Sigma$ iff Σ is refutable.*

Proof. For the direction from left to right, we will in fact show something stronger than the statement in the Proposition: we will prove the existence of a *single* model that refutes *every* sequent in Builder's winning region in $\mathcal{D}_{\mathbf{G}}@\Sigma$. The domain of this canonical countermodel consists precisely of these sequents themselves:

$$S := \mathsf{Seq}(\mathrm{ML}_\Lambda) \cap \mathsf{Win}_B(\mathcal{D}_{\mathbf{G}}).$$

For the definition of the coalgebra map and the valuation of this model we fix some *positional winning strategy* f for Builder in $\mathcal{D}_{\mathbf{G}}$. Recall that such a strategy exists by Theorem 8.12, and that it is a function that maps any position of the form $(\Theta, \mathsf{d}, \tau)$ to a premise of the rule instance that the position represents. Furthermore, we will confine attention to f-guided matches where Prover plays 'supportively'. To explain this we need some preparations.

First, we call a sequent Θ *propositionally simple* if each of its members has propositional complexity zero or one. In other words, a sequent is propositionally simple if it contains no formulas of the form $\neg\neg\varphi$, $\varphi \wedge \psi$ or $\neg(\varphi \wedge \psi)$; that is, all of its elements are atomic or modal formulas, or negations of such formulas. Note that only modal rules may be applied to a propositionally simple sequent (unless it is axiomatic).

For the definition of the coalgebra map $\sigma_f : S \to TS$ and the valuation V_f we confine our attention to f-guided matches where Prover plays 'supportively'. That is, we fix some strict linear order $\sqsubset$ on formulas that is compatible with the subformula order ($\varphi \triangleleft \psi$ implies $\varphi \sqsubset \psi$), and require that

- Prover postpones playing modal rules until no propositional rules are available; and
- when playing a propositional rule, Prover picks the principal formula that is *maximal* with respect to $\sqsubset$.

T **Claim 1** Let Θ be a sequent in S. If Prover plays supportively, then there is a *unique* f-guided partial match π such that $\mathsf{first}(\pi) = \Theta$, $\mathsf{last}(\pi)$ is propositionally simple, and P only plays propositional rules during π.

Given a sequent state $\Theta \in S$, we let π_Θ denote the match given by the above Claim, and we define

$$\begin{array}{lcl} \Theta^\ell & := & \mathsf{last}(\pi_\Theta) \\ \Theta^+ & := & \bigcup\{\Phi \mid \Phi \text{ occurs on } \pi_\Theta\}. \end{array}$$

In words, Θ^+ consists of all formulas that belong to some sequent on the path π_Θ.

T **Claim 2** For any $\Theta \in S$, the set Θ^+ is saturated; that is, for all formulas φ, ψ we have:
(1) it is not the case that $\varphi, \neg\varphi \in \Theta^+$ or that $\top \in \Theta$;
(2) if $\neg\neg\varphi \in \Theta^+$, then also $\varphi \in \Theta^+$;
(3) if $\varphi \wedge \psi \in \Theta^+$, then $\varphi \in \Theta^+$ or $\psi \in \Theta^+$;
(4) if $\neg(\varphi \wedge \psi) \in \Theta^+$, then also $\neg\varphi, \neg\psi \in \Theta^+$.
In addition we have that
(5) $\Theta^\ell \subseteq \Theta^+$;
(6) $\varphi \in \Theta^\ell \iff \varphi \in \Theta^+$, for all φ of the form $\pm p$ or $\pm\heartsuit_\lambda(\psi_0, \ldots, \psi_{n-1})$.

The definition of the valuation V_f is straightforward:

$$V_f(p) := \{\Theta \in S \mid \neg p \in \Theta^+\}.$$

For the coalgebra map σ_f we have to do more work; we will define σ_f layer by layer, where the layers of S are given by the modal depth of sequents. Recall the notion of the *modal depth* $\mathsf{md}(\Theta)$ of a sequent Θ from Definition 8.9.

Note that the sequents that we subsequently encounter in a match of the derivability game have *decreasing* modal depth. In particular, if Prover plays a modal rule at a sequent position Θ, then the next sequent of the match will have a *strictly smaller* modal depth than Θ: the modal rule peels off one layer of modalities. If she plays a propositional rule, the modal depth may decrease or stay the same, but it will not increase.

The layers of S are then given by the following definition:

$$S_n := \{\Theta \in S \mid \mathsf{md}(\Theta) \leq n\}.$$

By standardness of T have $TS_m \subseteq TS_n$ if $m < n$. The following observation is the pivotal claim in the proof.

Claim 3 Assume that, for some $n \in \omega$, there is a map $\sigma_n : S_n \to TS_n$ such that in the resulting coalgebra model $\mathbb{S}_n = (S_n, \sigma_n, V_f \restriction_{S_n})$ every state $\Phi \in S_n$ refutes the sequent Φ^+. Then we can extend σ_n to a map $\sigma_{n+1} : S_{n+1} \to TS_n$ satisfying the following two conditions:

(1) For every $\Theta \in S_{n+1} \setminus S_n$ and every formula $\heartsuit_\lambda(\psi_0, \ldots, \psi_{n-1})$ we have, writing $\tau := \sigma_{n+1}(\Theta)$:

$$\begin{array}{rll} \heartsuit_\lambda(\psi_0, \ldots, \psi_{n-1}) \in \Theta^+ & \text{implies} & \tau \notin \lambda_{S_n}([\![\psi_0]\!]^{\mathbb{S}_n}, \ldots, [\![\psi_{n-1}]\!]^{\mathbb{S}_n}) \\ \neg\heartsuit_\lambda(\psi_0, \ldots, \psi_{n-1}) \in \Theta^+ & \text{implies} & \tau \in \lambda_{S_n}([\![\psi_0]\!]^{\mathbb{S}_n}, \ldots, [\![\psi_{n-1}]\!]^{\mathbb{S}_n}). \end{array} \tag{69}$$

(2) In the resulting coalgebra model $\mathbb{S}_{n+1} = (S_{n+1}, \sigma_{n+1}, V_f \restriction_{S_{n+1}})$ every state $\Phi \in S_{n+1}$ refutes the sequent Φ^+.

Proof of Claim For part 1, fix some $\Theta \in S_{n+1} \setminus S_n$, and assume for contradiction that there is *no* object $\tau \in TS_n$ meeting the conditions in (69). To avoid notational clutter we assume that all predicate liftings in Λ are unary – this is not without loss of generality, but the general case is not essentially harder. Introduce a variable a_ψ for every formula ψ such that either $\heartsuit_\lambda\psi \in \Theta^+$ or $\neg\heartsuit_\lambda\psi \in \Theta^+$, for some $\lambda \in \Lambda$; let A be the set of all these variables; and let $\nu : A \to \mathrm{ML}_\Lambda$ be the natural substitution given by $v(a_\psi) := \psi$. Furthermore, define

$$\Gamma := \{\heartsuit_\lambda a_\psi \mid \heartsuit_\lambda\psi \in \Theta^+\} \cup \{\neg\heartsuit_\lambda a_\psi \mid \neg\heartsuit_\lambda\psi \in \Theta^+\},$$

then we obviously have $\Gamma[\nu] \subseteq \Theta^+$.

By our assumption on the non-existence of an element $\tau \in TS_n$ satisfying (69) we can find, for every $\tau \in TS_n$, a formula $\varphi_\tau \in \Theta^+$ such that either $\varphi_\tau = \heartsuit_\lambda\psi$ and $\tau \in \lambda_{S_n}([\![\psi]\!]^{\mathbb{S}_n})$ or $\varphi_\tau = \neg\heartsuit_\lambda\psi$ and $\tau \notin \lambda_{S_n}([\![\psi]\!]^{\mathbb{S}_n})$. Hence, if we define $m : S_n \to PA$ by

$$m : \Phi \mapsto \{a_\psi \in A \mid \mathbb{S}_n, \Phi \Vdash \psi\},$$

we obtain, for each $\tau \in TS_n$, a one-step formula α in Γ of the form $\alpha = \heartsuit_\lambda\psi$ (if $\heartsuit_\lambda\psi \in \Theta^+$) or $\alpha = \neg\heartsuit_\lambda a_\psi$ (if $\neg\heartsuit_\lambda\psi \in \Theta^+$) such that $\tau \in [\![\alpha]\!]^1_m$. In other words we find that $[\![\Gamma]\!]^1_m = TS_n$.

But then by one-step completeness there must be some modal rule $\mathsf{d} = \Delta_1 \cdots \Delta_{n-1}/\Gamma_0$ in $\mathbf{G}$ and a renaming $\tau : A \to A$ such that $\Gamma_0[\tau] \subseteq \Gamma$ and $[\![\Delta_i[\tau]]\!]^0_m = S$, for all i.

On the other hand, by construction of Γ we have $\Gamma_0[\nu \circ \tau] = \Gamma_0[\tau][\nu] \subseteq \Gamma[\nu] \subseteq \Theta$. But then by the observation in Claim 2(6) we have $\Gamma_0[\nu \circ \tau] \subseteq \Theta^\ell$, so that the triple $(\Gamma_0, \mathsf{d}, \nu \circ \tau)$ is a legitimate move for Prover at position Θ^ℓ. Then for each i, the sequent $\Delta_i[\nu \circ \tau]$ is a premise of this rule instance, and so one of these, say, $\Phi := \Delta_i[\nu \circ \tau]$ is the response picked by Builder's winning strategy f if Prover makes this move $(\Gamma_0, \mathsf{d}, \nu \circ \tau)$ indeed. Note that $\Phi = \Delta_i[\tau][\nu]$ implies

$$\Delta_i = \{a_\varphi \mid \varphi \in \Phi\} \cup \{\neg a_\varphi \mid \neg\varphi \in \Phi\},$$

and that $\mathsf{md}(\Phi) = \mathsf{md}(\Theta) - 1 = n$.

In other words, we have $\Phi \in S_n$, and so Φ, as a state in $\mathbb{S}_n$, refutes Φ^+ by assumption. That is, we find that $\mathbb{S}_n, \Phi \Vdash \varphi$ for no $\varphi \in \Phi^+$, which by definition of m means that $\Phi \notin [\![\Delta_i[\tau]]\!]^0_m$. This, however, contradicts the fact that $[\![\Delta_i[\tau]]\!]^0_m = S$.

For part 2 of the claim take an arbitrary sequent $\Theta \in S_{n+1}$, and distinguish cases. If Θ belongs to S_n, note that since σ_{n+1} extends σ_n we have $\sigma_{n+1}(\Theta) = \sigma_n(\Theta) \in TS_n$. In other words, (S_n, σ_n) is a subcoalgebra of (S_{n+1}, σ_{n+1}) and so we have that $\mathbb{S}_{n+1}, \Theta \Vdash \varphi$ iff $\mathbb{S}_n, \Theta \Vdash \varphi$, for *any* formula φ. In particular this means that Θ, as a state in $\mathbb{S}_{n+1}$, refutes Θ^+, by the assumption that it does so as a state in $\mathbb{S}_n$.

In case $\Theta \in S_{n+1} \setminus S_n$ it suffices to show that any formula $\varphi \in \mathrm{PL}(\Lambda(\mathtt{Sfor}_0(\Theta^+)))$ satisfies the following conditions:

$$\begin{array}{rcl} \varphi \in \Theta^+ & \text{implies} & \mathbb{S}_{n+1}, \Theta \not\Vdash \varphi \\ \neg\varphi \in \Theta^+ & \text{implies} & \mathbb{S}_{n+1}, \Theta \Vdash \varphi. \end{array}$$

We prove this statement by induction on the propositional complexity of φ.

In the base case of this induction φ is either atomic or of the form $\varphi = \heartsuit_\lambda(\vec{\psi})$. We leave the first case as an exercise for the reader. In the second case the above statements readily follow from part 1 of the Claim and the definition of $\mathbb{S}_{n+1}$.

In the induction step, we have two cases to consider. We leave the case where φ is of the form $\varphi = \neg\varphi'$ as an exercise for the reader. In the case where φ is of the form $\varphi = \varphi_0 \wedge \varphi_1$ we reason as follows. If $\varphi \in \Theta^+$ then by Claim 2 we find φ_i in Θ^+ for some $i \in \{0, 1\}$, so that by the induction hypothesis we obtain $\mathbb{S}_{n+1}, \Phi \not\Vdash \varphi_i$ for this i. But then obviously we do not have $\mathbb{S}_{n+1}, \Phi \Vdash \varphi$ either. On the other hand, if $\neg\varphi \in \Theta^+$ then by Claim 2 we find both formulas $\neg\varphi_0, \neg\varphi_1$ in Θ^+. By the induction hypothesis this means that $\mathbb{S}, \Theta \Vdash \varphi_i$ for both $i = 0, 1$, and clearly this implies that $\mathbb{S}, \Theta \Vdash \varphi_0 \wedge \varphi_1$, as required. $\lhd$

Based on Claim 3 we can now easily prove the following.

Claim 4 For every $n \in \omega$ there is a map $\sigma_n : S_n \to TS_n$ such that, in the resulting coalgebra model $\mathbb{S}_n = (S_n, \sigma_n, V_f)$ we have that Φ refutes Φ^+, for all $\Phi \in S_n$.

Proof of Claim We prove the claim by a straightforward induction. For the base case of this induction, where $n = 0$, observe that sequents in S_0 only contain propositional formulas. Given such a sequent Θ, it easily follows from Claim 2 that both Θ and Θ^+ are satisfiable sets of propositional formulas, and that, by the definition of the valuation V_f, Θ refutes Θ^+. However, we still need to define a coalgebra map $\sigma_0 : S_0 \to TS_0$, but for $\Theta \in S_0$ we may define $\sigma_0(\Theta)$ to be any member[12] of the set $T\{\Theta\}$. (Note that the definition of $\sigma_0(\Theta)$ has no impact on the truth of propositional formulas at Θ.)

In the inductive case we assume the existence of a map $\sigma_n : S_n \to TS_n$ satisfying the conditions of the claim, and we simply use the previous claim to find an appropriate map $\sigma_{n+1} : S_{n+1} \to TS_{n+1}$. $\lhd$

[12] In case this set is empty we are dealing with the constant functor $K_\varnothing$ which maps every set S to the empty set. This functor has only one coalgebra, namely, the empty one $(\varnothing, \varnothing)$. It follows that in this case *every sequent* is valid.

It is then a straightforward consequence of Claim 4 that the initial sequent Σ is refuted in the model $\mathbb{S}$. This finishes the proof of the left-to-right direction of the equivalence stated in the Proposition.

The opposite, right-to-left direction of the proof is left as an exercise for the reader.

QED

8.2 Hilbert calculi

▶ `Section to be added later`