An Abstract Look at the Fixed-Point Theorem for Provability Logic

Johan van Benthem

University of Amsterdam, Stanford University, and Tsinghua University

Abstract. In this brief note, I discuss some general backgrounds of the well-known Fixed-Point Theorem for Provability Logic, taking my cues from an elegant abstract argument by Dick de Jongh in the 1980s.

Keywords: modal provability logic · well-founded orderings · definability.

1 Introduction

In the mid 1980s, I was into 'semantic automata', (van Benthem, 1986), classifying linguistic quantifiers in terms of the complexity of their verification procedures on Venn diagrams. The next step in developing this 'procedural semantics' was an analysis of linguistic expressions that depend on the underlying structure of the object domain, and so, I developed an interest in tree automata whose computation rule is recursive in a given tree ordering. This led to an intensive and fruitful correspondence with Dick de Jongh about connections with provability logic, where such recursive definitions can be made explicit. In this correspondence, Dick came up with an elegant generalization of the key step in the Fixed-Point Theorem which applied far beyond the modalities, namely, to arbitrary generalized quantifiers satisfying suitable abstract conditions. Dick's result was included in my somewhat long and meandering paper 'Toward a Computational Semantics', (van Benthem, 1987) where it remained hidden.

The purpose of this brief note is twofold. I want to advertize Dick's result by itself, and its elegant level of abstraction. After that, I go further in this spirit and add some simple observations showing how the Fixed-Point Theorem can be seen as an instance of a family of abstract results on generalized well-founded orders. Much of what follows may be present in the folklore or the expert literature (more on this in Section 3), but a compact story may be useful.[1]

2 The Fixed-Point Theorem for Provability Logic

Consider the language of basic modal logic interpreted over relational models where the accessibility relation is transitive and upward well-founded. The standard modalities $\Box, \Diamond$ then range over all relational successors in this ordering.

[1] This paper offers a perspective that is complementary to (Litak & Visser, this volume), though the latter paper goes into far more proof-theoretic depth and width, and also covers different fixed-point constructions arising with binary modalities.

Call a modal formula $\varphi = \varphi(p, \mathbf{q})$ *p-modalized* if all occurrences of the proposition letter p in φ are in the scope of at least one modal operator $\Box$, while there is no syntactic constraint on the occurrences of the proposition letters in the sequence $\mathbf{q}$. We state the following result in its semantic version; a proof-theoretic version can be stated using the completeness theorem for Gödel-Löb logic.

Theorem 1 (Fixed-Point Theorem, De Jongh-Sambin-Bernardi). *For any p-modalized formula $\varphi(p, \mathbf{q})$, there exists a formula $\alpha = \alpha(\mathbf{q})$ such that $\varphi(\alpha(\mathbf{q}), \mathbf{q}) \leftrightarrow \alpha(\mathbf{q})$ is valid. This fixed-point formula α can be found by an effective procedure. Moreover, the fixed-point is unique up to logical equivalence.*

There are many proofs of the existence of fixed-points in provability logic, both model-theoretic and proof-theoretic, including a pleasantly non-constructive one via Beth's Theorem, while Maksimova also derived Beth's theorem from the fixed-point theorem, (Hoogland, 2001). A compact informative survey of provability logic can be found in (Verbrugge, 2017). The analysis of fixed-points in this paper is also indebted to a perspicuous proof in (Reidhaar-Olson, 1989).

Two major features. All these proofs make recurrent use of two general features. One is the semantic import of the restriction to p-modalized formulas $\varphi(p)$. In any model $\mathbf{M}$, their truth value at a point s is only sensitive to the intersection of the valuation set $V(p)$ of the model with the set of successors of s in $\mathbf{M}$. In *generalized quantifier* terminology [see below], the set $\{t \mid Rst\}$ is a *live-on set* for the quantifier defined by $\varphi(p)$ at s. The other feature is the availability of *induction principles* over well-founded orderings, in the form of Löb's Axiom or suitable variants. In fact, well-foundedness just *is* an induction principle.

The Recursion Theorem in set theory. A common feeling is that something general is going on in the above Fixed-Point Theorem beyond the specifics of a modal language. To see this, recall the Recursion Theorem from your introductory course in set theory, (Hrbacek & Jech, 1989), or perhaps you encountered an arithmetical version in recursion theory. One of its many versions says that functions defined by recursion on the transitive closure of their arguments in the epsilon relation exist and are unique. The proof shows striking similarities with the provability logic case: e.g., one first proves uniqueness by epsilon-induction, and then uses this to prove existence, again by epsilon induction. But the content of the Recursion Theorem is more sweeping: the language is not a simple modal one, but that of set theory, and also, the well-founded relation is not transitive.

One can reproduce the proof of the Recursion Theorem in more restricted formal languages. Still close to set theory is *monadic second-order logic* over well-founded orderings, but in what follows, we will also think of the system SOML: *second-order modal logic* with quantification over proposition letters. Fixed-points arguments in the setting of SOML specialize to simpler modal languages, and yet bring out the essence of proofs in richer formalisms.

But this is not yet the whole story. The striking fact about the Fixed-Point Theorem for provability logic is that the fixed-points exist *inside the first-order modal language* used to state the fixed-point equation. This brings us to the definability of fixed-points, and to De Jongh's style of analysis.

3 De Jongh's generalized quantifier analysis

A parametrized generalized quantifier (henceforth called a 'quantifier', for short) is a relation between points and sets of points. In any model $\mathbf{M}$, each modal formula $\varphi(p)$ defines a modal generalized quantifier $Q_{\varphi,p}$ as follows:

$$Q_{\varphi,p}sX \text{ iff } \mathbf{M}[p := X], s \models \varphi(p)$$

Alternatively, one can think of these point-dependent generalized quantifiers as functions from sets to sets, as with the next-approximation maps used in the semantics of the modal μ-calculus, (Venema, 2020). We now turn to two special properties for such quantifiers that will be central to what follows.

As is well-known, all modal formulas are *invariant for generated submodels*, (Blackburn, de Rijke & Venema, 2001), and thus, the generalized quantifiers $Q_{\varphi,p}sX$ defined in the above way only depend on the intersection of the set X with the set $\{s\} \cup \{t \mid Rst\}$. When $\varphi(p)$ is p-modalized, this strengthens to dependence only on the set of accessible points $X \cap \{t \mid Rst\}$. Let us call a quantifier satisfying the latter property *future-oriented*.

Next, here is a special property that is crucial to getting explicit fixed-points. We call a quantifier Q *persistent* if QsX implies QtX for all t with Rst: in modal terms, for persistent quantifiers, $Q\varphi \to \Box Q\varphi$ is valid.

With all this in place, here is Dick's result.

Fact 1 *All future-oriented persistent generalized quantifiers Q have the fixed-point $Q\top$, i.e., the equivalence $Q\top \leftrightarrow QQ\top$ is valid on well-founded models.*

We committed some harmless abuse of notation here, using 'Q' both in syntax and semantics. The proof to follow is the essence of a first step in a standard proof of the Fixed-Point Theorem for special formulas $\Box\varphi(p)$, (Smoryński, 1984).

Proof. We reason inside models $\mathbf{M}, s$ at some specific point s.

(i) $Q\top \to QQ\top$. Let $\mathbf{M}, s \models Q\top$. By persistence, we also have $\Box Q\top$, and therefore $\Box(Q\top \leftrightarrow \top)$. Using future-orientation, we get that $\mathbf{M}, s \models QQ\top$.

(ii) $QQ\top \to Q\top$ is proved using well-foundedness. Assume that $\mathbf{M}, s \models \neg Q\top$, then $\neg Q\top \wedge \Box Q\top$ is true at s or at some t with Rst. Using future-orientation as before, $\neg QQ\top$ is then true at t, and hence it is also true at s: either directly, or using persistence once more. (Note: The contraposition style in this second leg is just for convenience; one can also run it positively via well-founded induction.)

Remark on proof style. A word about this semantic style of reasoning, that will be used throughout. Uniqueness and existence hold locally at points s in models based only on what is true in the generated submodel at s. In a proof-theoretic version of our semantic arguments, this additional information would translate into explicit modal assumptions such as $\Box(p \leftrightarrow q)$ that keep track of exactly what is used. The proofs to follow should be understood with this in mind.

Fact 1 yields just one fixed-point, perhaps a trivial-looking one. But in fact, it is the only fixed-point up to logical equivalence. A standard argument for the following fact depends only on well-foundedness, and does not need persistence.

Fact 2 *If a future-oriented quantifier has a fixed-point, that fixed-point is unique.*

What Dick's analysis shows is how a well-chosen notation identifies abstract assumptions that go into more complex syntactic arguments. One is reminded of how provability logic itself clarifies key patterns in complex metamathematical proofs with well-chosen simple notation. There is also the beauty of abstraction, and the suggestion of more general perspectives.

In particular, there is nothing sacrosanct about the basic modal setting. The above result also applies, say, to *graded modal logic*, since the modality

$[\leq \mathrm{k}]\varphi$ 'φ is true in all except k successors'

is future-oriented and persistent. As an example of a de Jongh fixed-point in graded modal logic, consider the formula $p \leftrightarrow [\leq 1]p$, which might be read as 'up to at most one exception, p is refutable'. The fixed-point $[\leq 1]\bot$ is only true in endpoints or points that see only one endpoint. But in fact, Dick's result applies to any quantifier $\Box Q$, whether Q is modal, first-order, or not: say,

in *most* successors of each successor, or

for each successor, in an *even number* of its successors

In finite at most binary trees, the de Jongh fixed-point for the latter quantifier defines the 'well-balanced trees' where no node has just one successor.

Languages with such extended generalized quantifiers go far beyond the original arithmetical motivations for provability logic, and have their uses elsewhere. But still, one might be allowed to wonder whether, e.g., the two just-mentioned modalities correspond to weaker notions of arithmetical provability, say, allowing small sets of exceptions in a numerical or measure-theoretic sense.

Still, we are not done yet, as we want fixed-point theorems for complex formulas, not just for single generalized quantifiers. I believe this can be done in a suitably generalized quantifier perspective, but in this paper, I follow a more standard route, borrowing known ideas from provability logic.

Historical comment In response to a first draft of this paper, Albert Visser pointed me to (Smoryński, 1985) which has Dick's idea in a proof-theoretic setting. In Chapter 4, Smoryński introduces a theory SR with two basic axioms:

(a) $\triangle\varphi \to \Box\triangle\varphi$

(b) $\Box(\varphi \leftrightarrow \psi) \to (\triangle\varphi \leftrightarrow \triangle\psi)$,

where $\triangle$ is a modal operator with an arbitrary interpretation, for which the book gives concrete examples in the realm of arithmetical provability. Smoryński proves various results, and claims a full fixed-point theorem. Given this, it seems that the ideas reported here may well have independent double origins.

For now, we first generalize in another direction, showing that fixed-points exist just as well with not necessarily transitive well-founded relations. This will yield a general fixed-point definition of which Dick's $Q\top$ is a special case.

4 Fixed-points on generalized well-founded orders

Dropping the transitivity assumption on relational models makes sense in many settings, for instance, with tree automata for procedural semantics that work in layers, (van Benthem, 1987). Recursion in non-transitive settings also drives the fixed-point logic analysis of Zermelo's Theorem and Backward Induction solutions for games in (van Benthem & Gheerbrant, 2010).

In this setting, we need to assume *well-foundedness* in the following sense: there are no infinite upward sequences $s_1 R s_2 R s_3, \ldots$ in the ordering R (this terminology is for brevity: the modal literature usually calls this 'converse well-foundedness'). This still gives the following induction principle to work with: each non-empty set X has a maximal element $s \in X$ in the sense that no t with $R^* st$ is in X, where, in this paper, R^* denotes the transitive closure of R. Restated in positive terms, this is an abstract form of course-of-values induction.

We now also generalize the notion of *future-orientation* for quantifiers Qp to dependence on the intersection of $V(p)$ with $\{t \mid R^* st\}$. The following assertion is easy to prove by the new form of induction, in a way that resembles, as announced earlier, the proof of the Recursion Theorem in set theory.

Fact 3 *If Q is future-oriented, then it has a unique fixed-point.*

The following gives more information about the shape of this fixed-point.

Proof. (i) Uniqueness follows by a standard well-foundedness argument. Beyond a last point s where two putative fixed-points p, q differ, they coincide on the points reachable from s in the transitive closure R^*. But at that s, by future-orientation, Qp, Qq have the same truth value, and hence so do p, q. (ii) As for existence, the following second-order formula in fact defines the fixed-point:

$$\exists p\,(Qp \wedge [R^*](Qp \leftrightarrow p))$$

Call this formula α: we will prove that $Q\alpha \leftrightarrow \alpha$ by the above form of induction. Assume that $[R^*](Q\alpha \leftrightarrow \alpha)$ is true. We show that $Q\alpha \leftrightarrow \alpha$ is true as well.

(a) Let $Q\alpha$ be true. Then α describes a unary predicate (or set) p such that $Qp \wedge [R^*](Qp \leftrightarrow p)$, and the latter fact means by definition that α holds.

(b) Let α be true. That is, for some p we have that $Qp \wedge [R^*](Qp \leftrightarrow p)$ is true. We also have $[R^*](Q\alpha \leftrightarrow \alpha)$, and so, by unicity of fixed-points, we have $[R^*](\alpha \leftrightarrow p)$. Using future-orientation w.r.t. Qp, we then get $Q\alpha$.

Now Dick's result follows as a consequence of this general analysis.

Corollary 1. *The preceding result implies Fact 1.*

Proof. We show that, if the quantifier Q is persistent, then the above formula $\alpha = \exists p\,(Qp \wedge [R^*](Qp \leftrightarrow p))$ is in fact equivalent with $Q\top$.

(i) Suppose that $Q\top$ holds at s, then it holds in the whole generated submodel at s, by repeated applications of persistence. Therefore, $Q\top$ is equivalent to $\top$ in the whole generated submodel, and α is true by taking p to be $\top$.

(ii) Now let α hold at s. Then, for some p, Qp is true, and so by persistence, Qp is equivalent to $\top$ in all reachable successors. Also, we have $[R^*](Qp \leftrightarrow p)$, and so $[R^*](\top \leftrightarrow p)$. By the future-orientation of Qp, $Q\top$ then holds at s.

5 Fixed-point theorems for generalized modal languages

Next, consider a propositional modal language with many different modalities [first-order or higher-order], viewed as generalized quantifiers $Q_1, \ldots, Q_k$ that are all future-oriented in the sense of the preceding section. Fixed-point equations are formulas $p \leftrightarrow \psi(p)$ for possibly complex p-modalized formulas ψ in this language. Each such formula can be viewed as being of the form

$$\varphi[Q_1\alpha_1(p), \ldots, Q_k\alpha_k(p)]$$

where the displayed subformulas list all occurrences of p in the p-free 'skeleton formula' φ. (Incidentally, while such decompositions always exist, they need not be unique.) Clearly, such a formula can also be seen as defining one generalized quantifier. w.r.t. the propositional variable p. Moreover, given the invariance for generated submodels and the future-orientation of all the quantifiers occurring inside the formula, this overall quantifier is future-oriented.

Let us spell out this fact, as it also underlies later proofs. Let Q be future-oriented and α invariant for generated submodels. Consider a point s in a model $\mathbf{M}$ where (i) $[R^*](p \leftrightarrow q)$ and (ii) $Q\alpha(p)$ are true. By (ii), $Qs\,\{t \mid \mathbf{M} \models \alpha(p)\}$ which holds iff $Qs\,\{t \mid \mathbf{M_s} \models \alpha(p)\}$, with $\mathbf{M}_s$ the generated submodel of $\mathbf{M}$ at point s. Next, by the future-orientation of Q, (i) implies that $Qs\,\{t \mid \mathbf{M_s} \models \alpha(q)\}$, and by invariance for generated submodels once more, we get $Qs\,\{t \mid \mathbf{M} \models \alpha(q)\}$. [Actually, less is needed than invariance for generated submodels to make this argument work, but we are not aiming for a minimal version in this paper.]

Now, given the future-orientation, by Fact 3, there is a unique fixed-point for $\psi(p)$ which can be defined explicitly by one unary existential second-order quantifier $\exists p$ over a formula of the given quantifier language (at least, if we also have a transitive-closure modality for the well-founded order in the language). This analysis leads to the following result for second-order modal languages.

Fact 4 *The Fixed-Point Theorem holds for second-order modal logic over a well-founded order with arbitrary added future-oriented modalities.*

Proof. The above analysis applies to the extended modal language of SOML, where in particular, each formula is still invariant for generated submodels. Therefore, our abstract fixed-point of Fact 3 is available. Moreover, we need no separate modality for the operator $[R^*]$ in the definition, since this is known to be definable as a unary fixed-point which can itself be defined in SOML.

An alternative analysis of Fact 3 shows that the Fixed-Point Theorem holds for extended modal logics that allow for elimination of 'bisimulation quantifiers', (Hollenberg, 1996), such as suitable fragments of the modal μ-calculus.

Now let us return to basic modal languages without second-order quantifiers. Even here, lifting the transitivity constraint is not unnatural. (Van Benthem, 2006) studies variants of Löb's Axiom using transitive closure modalities which separate out its well-foundedness content from its transitivity content.

The following extension of Dick's result generalizes the Fixed-Point Theorem for provability logic, but cf. our remark about (Smoryński, 1985).

Theorem 2. *For any modal language over a well-founded order whose formulas are invariant for generated submodels and whose vocabulary contains arbitrary future-oriented persistent modalities, each fixed-point equation has a unique solution definable inside the language.*

Proof. We follow the proof in (Reidhaar-Olson, 1989), but cutting out the proof-theoretic excursions, and taking a direct route to existence. Moreover, unlike in that proof, we use the earlier abstract solution of Fact 3 at a crucial stage.

Also as before, we reason locally at some arbitrary point s in a model. Consider a formula $\psi(p) = \varphi[Q_1\alpha_1(p), \ldots, Q_k\alpha_k(p)]$ as described above. This describes a unique fixed-point α with a second-order definition as stated earlier in the proof of Fact 3, so the following formula is true throughout our model:

$$\alpha \leftrightarrow \varphi[Q_1\alpha_1(\alpha), \ldots, Q_k\alpha_k(\alpha)]$$

Here, since combinations $Q_i\alpha_i(p)$ with formulas α_i that are themselves invariant under generated submodels still define future-oriented invariant quantifiers, we can simplify notation in what follows without loss of generality to

$$\psi(p) = \varphi[Q_1 p, \ldots, Q_k p]$$

where the unique fixed-point simplifies to:

$$\alpha = \varphi[Q_1(\alpha), \ldots, Q_k(\alpha)]$$

Now we prove the definability of this fixed-point inside our language. The base case with $k = 0$ is obvious, as the formula ψ itself is then a fixed-point.

Next, let $k > 0$, and assume that we have the result already for all formulas with fewer than k occurrences of p in the display. Define ψ_i as the formula

$$\varphi[Q_1 p, \ldots, \top, \ldots, Q_k p]$$

where the i-th occurrence $Q_i p$ has been replaced by $\top$. By the inductive assumptions, this has a fixed-point δ_i satisfying

$$\delta_i \leftrightarrow \varphi[Q_1(\delta_i), \ldots, \top, \ldots, Q_k(\delta_i)]$$

Lemma 1. *The following equivalence holds everywhere: $Q_i(\delta_i) \leftrightarrow Q_i(\alpha)$.*

Proof. (i) First assume that $Q_i(\delta_i)$ is true at a point s. By the persistence of Q_i, we have as in earlier arguments that $(Q_i(\delta_i) \leftrightarrow \top) \wedge [R^*](Q_i(\delta_i) \leftrightarrow \top)$ is true at s. By invariance of φ for generated submodels and replacement of equivalents in the generated submodel at s, this implies that we have the following equivalence true throughout this generated submodel:

$$\delta_i \leftrightarrow \varphi[Q_1(\delta_i), \ldots, Q_i(\delta_i), \ldots, Q_k(\delta_i)]$$

But then δ_i is a fixed-point for the original formula ψ, and as fixed-points are unique, we have $\delta_i \leftrightarrow \alpha$ true in the generated submodel, with α the earlier fixed-point. It then follows from the truth of $Q_i(\delta_i)$ that $Q_i(\alpha)$ is true as well.

(ii) Next let $\neg Q_i(\delta_i)$ be true at a point s. Then by well-foundedness, there is a point t R-reachable from s in 0 or more steps with $\neg Q_i(\delta_i) \wedge [R^*]Q_i(\delta_i)$ true. In this point, we also have $[R^*](Q_i(\delta_i) \leftrightarrow \top)$ true. As before, it follows that

$$[R^*](\delta_i \leftrightarrow \varphi[Q_1(\delta_i), \ldots, Q_i(\delta_i), \ldots, Q_k(\delta_i)])$$

and hence, by uniqueness of fixed-points again, $[R^*](\delta_i \leftrightarrow \alpha)$ is true at t. Then, by the future-orientation of Q_i, $\neg Q_i(\delta_i)$ implies that $\neg Q_i(\alpha)$ is true at t: and, either directly, or using persistence of Q_i as often as necessary on the finite path from s to t, we have that $\neg Q_i(\alpha)$ is true at s.

The rest of the argument is now easy. Using the lemma repeatedly, together with semantic replacement of equivalents, we can replace the successive occurrences of α in $\varphi[Q_1(\alpha), \ldots, Q_k(\alpha)]$ to obtain the equivalence

$$\alpha \leftrightarrow \varphi[Q_1(\delta_1), \ldots, Q_k(\delta_k)]$$

where the formula to the right-hand side is inside our modal language.

I believe that the preceding proof identifies the essentials of the Fixed-Point Theorem in its original setting. But in line with what was said earlier about Dick's result, it applies much more widely, to graded modal logic, modal logics with persistent future-oriented second-order quantifiers, and so on. As a simple example, it is easy to compute a counterpart for the Löb fixed-point in terms of graded 'almost provability': $[\leq 1]\neg p \rightarrow q$ has the fixed-point $[\leq 1]\neg q \rightarrow q$. More exciting examples can be found with iterated persistent graded modalities.

All this expressive power still leaves an issue that may hamper the generality. Normally, in extended modal languages, modalities can refer to *many relations*, whereas here, all modalities refer to one fixed well-founded order R. Can this be lifted? A few thoughts on this will be found in Section 7.

6 Fixed-point theorems for non-persistent modalities

Persistence is a sufficient condition on future-oriented quantifiers to guarantee a fixed-point theorem. But is it necessary? As reported in (van Benthem, 1987), De Jongh also had a nice example showing that things at least become more difficult without persistence. Consider the simple fixed-point equation

$$p \leftrightarrow \mathrm{MOST}\neg p$$

where 'MOST' is the future-oriented but not persistent modality 'in most successors'. When computing this fixed-point on the natural numbers $\{0, 1, 2, \ldots\}$

with $>$, p turns out to be true in all *odd numbers*. But the property of being odd is not definable in the pure modal language with MOST.

The fixed-point in the preceding example is still quite regular. Indeed, in the model of the preceding example, the odd positions are definable as follows in the language of *propositional dynamic logic* PDL over the immediate predecessor relation R_a. This widely used formalism shifts our earlier notation slightly, and in what follows, Kleene star will denote *reflexive-transitive closure*:

$$\langle a; (a; a)^* \rangle [a] \bot$$

Here is a general fact behind the preceding observation. Consider upward well-founded deterministic relations R_a, i.e., taking the transitive closure: finite linear orders, or in terms of formal languages and automata theory, *finite words*. Take a propositional modal language of '*bare* PDL' with propositional constants $\top, \bot$, one propositional variable p, Boolean operations, and modalities $\langle \pi \rangle$ where π is a program expression constructed from a using just the regular operations ; $\cup$, $*$. In this setting, a formula $\varphi(p)$ is called *p-modalized* if each occurrence of p in $\varphi(p)$ occurs in the scope of at least one modality $\langle a$. (This syntactic format can be generalized, since in the proof to follow, we could also allow modalities for complex PDL-programs that are equivalent to $a\,;\,\pi$ for some program π.)

Fact 5 *Fixed-point equations $p \leftrightarrow \varphi(p)$ with p-modalized $\varphi(p)$ over finite words in an alphabet $\{a\}$ are uniquely solvable with a fixed-point defined in bare* PDL.

Proof. Consider the natural numbers with the immediate predecessor relation R_a. Computing the fixed-point upward from the endpoint 0 yields a unique pattern of truth values for p, since $\varphi(p)$ does not refer to any structure except the bare order. In this structure, any modal formula defines a set of points, or equivalently, because of the invariance of modal formulas for generated submodels, a set of linear orders. Also, given the finite modal depth k of $\varphi(p)$, the truth value of p at a point only depends on the truth values of p in the k preceding positions. Call any occurring distribution of truth values for p over k successive positions a *p-segment*. We now analyze what this means combinatorially.

Looking at successive intervals of length k, there must be a p-segment S that occurs infinitely often, possibly separated by other intervals. Moreover, given the functional nature of the computation for $\varphi(p)$, the truth values for p to the right of occurrences of S are all in the same order. Therefore, there is in fact a unique interval T separating the occurrences of S. Thus, from some finite arbitrary interval of truth values for p onward, the whole ordering consist of repetitions of the interval $S \cup T$. But then we can define p as a disjunction of

> (a) Explicit modal definitions for the p-positions in the arbitrary initial interval, say, $\langle a\,;\,a \rangle \neg \langle a \rangle \top$ for distance 2 to the endpoint, (b) For each p-position in $S \cup T$: a modal formula describing the distance of that position to the end of $S \cup T$ followed by $\langle (a^m)^* \rangle \langle a^n \rangle \neg \langle a \rangle \top$, where m is the length of $S \cup T$, and n is the length of the arbitrary initial interval.

In terms of natural numbers, this defines *semi-linear sets* with just one period.

One would like to generalize to full PDL with more atoms and tests inside programs, over branching well-founded orders. However, this seems infeasible.

Counter-example: solving games. Zermelo's Theorem says that all finite-horizon two-player zero-sum games of perfect information are determined: one of the players i, j has a winning strategy. The key to its proof is the following recursive definition of the positions in the game tree where player i has a winning strategy. Here `move` stands for the union of all moves, $\mathrm{win_i}$ marks winning end positions for player i, and $\mathrm{turn_i}$ marks turns for i:

$$\mathrm{WIN_i} \leftrightarrow$$
$$((\neg\langle \mathtt{move}\rangle\top \wedge \mathrm{win_i}) \vee (\mathrm{turn_i} \wedge \langle \mathtt{move}\rangle \mathrm{WIN_i}) \vee (\mathrm{turn_j} \wedge [\mathtt{move}]\mathrm{WIN_i}))$$

This is a fixed-point equation of the form

$$p \leftrightarrow (q \vee (r \wedge \Diamond p) \vee (s \wedge \Box p))$$

which does not seem to yield an explicit definition for p inside the language of PDL. The intuitive reason is that the solution pattern $(\Diamond\Box)^* q$ over finite trees demands the existence of a subtree encoding the winning strategy, as opposed to a mere path, (van Benthem, 2006).

Still, at present, this is just a conjecture: it seems an open problem whether Zermelo-style fixed-points can be defined in PDL. Is propositional dynamic logic closed under fixed-points on finite trees, or on well-founded orders in general?

A simple fixed-point description. Consider a basic modal fixed-point formula $\varphi(p)$ with p modalized, on finite trees which has a fixed-point α by our general results. Since subformulas $\Box\psi$ are true in endpoints, evaluating $\varphi(p)$ on endpoints makes this formula equivalent with a Boolean combination of proposition letters other than p, say, φ^0. Next, for points at level 1 (where all successors are endpoints), we substitute φ^0 for p to obtain a formula $\varphi^1 = \varphi(\varphi^0/p)$ which is easily seen to be true at the same points at levels $0, 1$ as the fixed-point α. Working upward inductively with successive substitutions, we obtain formulas φ^k for each natural number k. Given this construction, it is easy to prove the following fact:

On points up to level k, φ^k is equivalent with the fixed-point formula α.

Let us define φ^* as the countable disjunction of all formulas φ^k. It follows from the preceding observation that φ^* defines the fixed-point of $\varphi(p)$.

Modal substitution logic. This suggests an extension of the basic modal language to a system MSL with a minimum needed for the preceding to work. We introduce explicit notation for *substitutions* $p := \psi$ that are interpreted semantically as transforming models $\mathbf{M}$ by resetting the valuation for p to the truth set $[[\psi]]^{\mathbf{M}}$. Moreover, we add *regular operations* over these transformations, in particular, iterations $(p := \psi)^*$. Thus, the syntax of MSL is that of PDL, now with formulas plus expressions for model-transforming programs, but it is more powerful.

For instance, the above pattern $(\Diamond\Box)^*\varphi$ is definable as

$$\langle p := \varphi \,;\, (p := \Diamond\Box p)^* \rangle \top$$

where p is a fresh proposition letter not occurring in φ. The above fixed-points are all definable in this formalism, using finitary iterations of substitutions.

The language of the system MSL may look unusual, but it is finitary, and it is semantically modal in that its formulas are invariant for bisimulation. One can also show that MSL extends PDL, while, on finite trees, it is contained in the modal μ-calculus. The general case may have higher complexity, like the logic of iterated relativization modalities in (Miller & Moss, 2005). A further investigation of modal substitution logic is left to another occasion.

The use of dynamic modalities for model transformation in MSL brings iterated manipulations into the syntax of a logic in a generic manner. An alternative format inspired by fixed-point analysis is the graph-based 'cyclic Henkin syntax' of (Visser, 2021), and it would be of interest to compare the two approaches.

7 Further directions

At this point, many directions open up. One can extend our analysis to multiple fixed-points for *simultaneous fixed-point equations* and to generalized quantifiers with more arguments. One can move from model theory to *proof theory*, and determine the precise content of the preceding semantic arguments and the logical strength they presuppose (classical, intuitionistic, or still less). Also, the generalized quantifier perspective employed here suggests analogies with *neighborhood semantics* for modal logic, (Pacuit, 2017), since the usual modal neighborhood frames are nothing but point-dependent generalized quantifiers in our sense. Instead of pursuing these issues, we mention three other points of interest.

Other fixed-point logics. Provability logic is one modal fixed-point logic, more widely used is the *modal μ-calculus.* The two systems are related: (van Benthem, 2006) embedded provability logic faithfully into the μ-calculus, (Visser, 2005) has a partial converse. Also, the mechanics of the two systems are close: computing fixed-points for modalized formulas on well-founded orders in μ-calculus approximation style eventually stabilizes to the provability logic fixed-points, despite negative occurrences of the running variable p. Taking our generalized quantifier perspective to the μ-calculus is natural, since that system also works for *arbitrary monotonic modalities*, not just standard ones or first-order quantifiers. One can even move the analysis to first-order fixed-point logic LFP(FO), (Ebbinghaus & Flum, 1995) or its fragment of 'guarded fixed-point logic', (Grädel, 2002).

Connections with automata theory and tree logics. The analysis in Section 6 employs essentially weak monadic second-order logic of one successor for finite strings, a special case of monadic second-order logic MSOL over strings which defines the regular languages, (Büchi, 1960). We believe that the classical results by Büchi and Rabin connecting MSOL over strings and trees with second-order languages, (Grädel, Thomas & Wilke, 2002), have a deeper bearing on the fixed-point theorems discussed in this paper. In fact, automata provide an alternative syntax format whose state graphs are more tolerant of defining fixed-points. However, the precise connection remains to be understood.

Extending the languages sideways. Modal languages are often about many orderings connected in certain ways. Also, general fixed-point logics can have a lot of vocabulary referring to complex relational similarity types for their models. Of course, there is a distinguished well-founded order R in the background: often the order of the ordinals, that can be added to models explicitly as a further component when proving basic results such as the Downward Löwenheim-Skolem theorem for $\mathsf{LFP}(\mathsf{FO})$, (Flum, 1995). What happens to the fixed-point theorem in the presence of other relations with matching modalities?

Some extensions are straightforward. For a formula $p \leftrightarrow (\langle S\rangle q \vee \neg[R]p)$, with R a well-founded relation and S an arbitrary further relation, a fixed-point can be computed as usual, treating the disjunct $\langle S\rangle q$ as a proposition letter. For more general cases, modalized occurrences of p must occur under modalities for relations that are well-founded (these can be basic, but combinations $R\,;S$ can be well-founded too). Also acceptable are combinations like $\langle R\cap S\rangle p$, as this just states R-reachability by a special route. If we do not require all this, fixed-points may still exist, but easy examples show that they need not be unique. (We could even allow occurrences of p that are not guarded in this manner as long as they are positive, and thus support a μ-calculus style recursion.)

A good test case for lifting the results of this paper to the broader polymodal settings suggested here might be the *Bounded Fragment* of first-order logic, where all formulas are invariant for generated submodels, (Feferman & Kreisel, 1966), (Areces, Blackburn & Marx, 2001).

8 Conclusion

The fixed-point theorem for provability logic and what makes it tick remains forever intriguing, and we owe this wide world to Dick de Jongh and his friends.

Acknowledgements I thank the referees for this volume, as well as Thomas Icard, Wolfgang Thomas, Yde Venema, Rineke Verbrugge, and especially, Albert Visser for their comments on this paper when still under construction.

References

1. C. Areces, P. Blackburn & M. Marx, Hybrid Logics: Characterization, Interpolation and Complexity, *Journal of Symbolic Logic*, 66:3, 977–1010, 2001.
2. J. van Benthem, *Essays in Logical Semantics*, Reidel, Dordrecht, 1986.
3. J. van Benthem, Towards a Computational Semantics. In P. Gärdenfors, ed., *Generalized Quantifiers: Linguistic and Logical Approaches*, Reidel, Dordrecht, 31–71, 1987.
4. J. van Benthem, Modal Frame Correspondences and Fixed–Points, *Studia Logica*, 83:1, 133–155, 2006.
5. J. van Benthem & A. Gheerbrant, Game Solution, Epistemic Dynamics, and Fixed–Point Logics, *Fundamenta Informaticae*, 100, 19–41, 2010.
6. P. Blackburn, M. de Rijke & Y. Venema, *Modal Logic*, Cambridge University Press, Cambridge: UK, 2001.

7. J. Büchi, Weak Second-Order Logic and Finite Automata, *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik*. 6, 66–92, 1960.
8. H-D Ebbinghaus & J. Flum, *Finite Model theory*, Springer Science Publishers, Heidelberg, 1995.
9. S. Feferman & G. Kreisel, Persistent and Invariant Formulas Relative to Theories of Higher Order, *Bulletin of the American Mathematical Society*, 72, 480–485, 1966.
10. J. Flum, On the (Infinite) Model Theory of Fixed-Point Languages, Mathematical Institute, University of Freiburg.
11. E. Grädel, Guarded Fixed-Point Logics and the Monadic Theory of Countable Trees, *Theoretical Computer Science*, 288, 129–152, 2002.
12. E. Grädel, W. Thomas & T. Wilke, eds., *Automata, Logics, and Infinite Games*, Lecture Notes in Computer Science, 2500, Springer, Berlin, 2002.
13. M. Hollenberg, 1996, *Logic and Bisimulation*, PhD Thesis, Department of Philosophy, University of Utrecht.
14. E. Hoogland, *Definability and Interpolation: Model-theoretic Investigations*, PhD Thesis DS-2001-05,, ILLC, University of Amsterdam, 2001.
15. K. Hrbacek & Th. Jech, *Introduction to Set Theory. Revised and Expanded*, CRS Press, Boca Raton, 1999.
16. T. Litak & A. Visser, Lewisian Fixed Points I:Two Incomparable Constructions, this volume, 2021.
17. J. Miller & L. Moss, The Undecidability of Iterated Modal Relativization, *Studia Logica*, 79, 373–4-7, 2005.
18. E. Pacuit, *Neighborhood Semantics for Modal Logic*, Springer Science Publishers, Beijing, 2017.
19. L. Reidhaar-Olson, A New Proof of the Fixed-point Theorem of Provability Logic, *Notre Dame Journal of Formal Logic*, 31:1, 37–43, 1989.
20. C. Smoryński, Modal Logic and Self-Reference. In D. Gabbay and F. Guenthner, eds., *Handbook of Philosophical Logic*, Reidel, Dordrecht, 441–495, 1984.
21. C. Smoryński, *Self-Reference and Modal Logic*, Springer Science Publishers, New York, 1985.
22. Y. Venema, *Lectures on the Modal μ-Calculus*, ILLC, University of Amsterdam, 2020.
23. R. Verbrugge, Provability Logic, *Stanford Encyclopedia of Philosophy*, Stanford, 2017.
24. A. Visser, Cyclic Henkin Logic, manuscript, Department of Philosophy, University of Utrecht, https://arxiv.org/abs/2101.11462, 2021.
25. A. Visser, Löb's Logic Meets the μ-Calculus, in *Processes, Terms and Cycles: Steps on the Road to Infinity*, Springer Lecture Notes in Computer Science, 3838, 14–25, 2005.
26. A. Visser, Provability Logic and Modalized Fixed Points, lecture slides, Online Seminar on Foundations of Mathematics, Wuhan University, 2020.