A New Fixed-Point Theorem for Logic Programming Semantics

Pascal Hitzler* and Anthony K. Seda
Department of Mathematics, University College Cork
Cork, Ireland
E-mail: {phitzler,aks}@ucc.ie
Web: http://maths.ucc.ie/~{pascal,seda}

Abstract

We present a new fixed-point theorem akin to the Banach contraction mapping theorem, but in the context of a novel notion of generalized metric space, and show how it can be applied to analyse the denotational semantics of certain logic programs. The theorem is obtained by generalizing a theorem of Priess-Crampe and Ribenboim, which grew out of applications within valuation theory, but is also inspired by a theorem of S.G. Matthews which grew out of applications to conventional programming language semantics. The class of programs to which we apply our theorem was defined previously by us in terms of operators using three-valued logics. However, the new treatment we provide here is short and intuitive, and provides further evidence that metric-like structures are an appropriate setting for the study of logic programming semantics.

Keywords: Logic Programming, Denotational Semantics, Supported Model, Generalized Metric, Fixed-point Theorem

Introduction

One advantage possessed by a logic program P, or a disjunctive database, over conventional imperative and object oriented programs is that it has a natural machine-independent meaning, namely, its logical meaning or declarative semantics. This is usually taken to be some "standard" model canonically associated with P. Unfortunately, there is often many possible choices for the standard model such as the well-founded model (van Gelder et al.), the stable model (Gelfond and Lifschitz), the perfect and weakly perfect models (Przymusinski), and so on, which do not in general coincide and all of which have a claim to be "the natural choice" depending on one's view of non-monotonic reasoning. It is therefore important and interesting to know when these various models coincide since this confirms coincidence of the various ways of considering non-monotonic reasoning.

In [6, 10], the authors defined certain classes of programs, called Φ-accessible and Φ^*-accessible programs, which have the property that each program in the class has a unique supported model, and showed that it follows from this property that all the different semantics mentioned above in fact coincide. These latter classes were defined in terms of various three-valued logics and are known to include the acceptable programs of Apt and Pedreschi (important in termination analysis), see [2], and certain other important classes, and are also known to be computationally adequate i.e. can compute all partial recursive functions; they therefore provide a semantically unambiguous setting with enhanced syntax and full computational power.

The supported models of P (or the Clark completion semantics) are the fixed points of the single-step operator T_P, and the proof of their existence and uniqueness we gave in [6] for the Φ-accessible and Φ^*-accessible programs was by means of three-valued logic. In this paper, we provide an alternative proof based on a new fixed-point theorem we establish here which generalizes the theorem of Priess-Crampe & Riben-

*The first named author acknowledges financial support under grant SC/98/621 from Enterprise Ireland.

boim [7, 14, 15]. This generalization was inspired by the occurrence of distance functions d with the slightly surprising property that $d(x, x)$ need not be zero. Indeed, such distance functions have already been discussed by Matthews in [12, 13] in connection with the data flow networks of Kahn, and this suggests that theorems of the sort we present here may have other interesting applications within computer science.

Preliminaries

A (*normal*) *logic program* is a finite set of universally quantified clauses, from first order logic, of the form

$$\forall\,(A \leftarrow L_1 \wedge \cdots \wedge L_n)\,,$$

where A is an atom and $L_1, \ldots, L_n$ are literals. Such clauses are usually written as

$$A \leftarrow L_1, \ldots, L_n.$$

We call A the *head* of the clause, and $L_1, \ldots, L_n$ the *body* of the clause. Each L_i is said to be a *body literal* of the clause. We refer to [11] for notation and basic concepts in logic programming.

For a given logic program P, we denote the Herbrand base (i.e. the set of all ground atoms in the underlying first order language) by B_P. As usual, (Herbrand-) interpretations of P can be identified with subsets of B_P, so that the set I_P of all interpretations of P can be identified with the power set of B_P. The set of all ground instances of each clause in a program P will be denoted by ground(P). A *level mapping* is a function $l : B_P \to \alpha$, where α is an arbitrary (countable) ordinal; we always assume that l has been extended to all literals by setting $l(\neg A) = l(A)$ for each $A \in B_P$.

The standard approach to logic programming semantics, i.e. to assigning a reasonable meaning to a given logic program, is to identify models of the program which have certain additional properties. We will focus here on the *supported model semantics* or *Clark completion semantics*, see [4, 1]. We define the *immediate consequence* or *single-step operator* T_P for a given logic program P as a mapping $T_P : I_P \to I_P$ of interpretations to interpretations as follows: $T_P(I)$ is the set of all $A \in B_P$ such that there exists a ground instance $A \leftarrow L_1, \ldots, L_n$ of a clause in P with head A and such that $I \models L_1 \wedge \cdots \wedge L_n$. Note that T_P is in general not monotonic. As it turns out, the models of P are exactly the pre-fixed points of T_P, that is, satisfy $T_P(I) \subseteq I$. A *supported model* (or model of the Clark completion [4]) of P is a fixed point of T_P, see [11] for these and related concepts.

The following definition is taken from [2] where it was employed in defining acceptable programs. Such programs have been shown to be of great importance in termination analysis, and we will use it as the basis of the more general Definition 2.

Definition 1 Let P be a logic program and let p, q be predicate symbols occurring in P.

1. p *refers to* q if there is a clause in P with p in its head and q in its body.
2. p *depends on* q if (p, q) is in the reflexive, transitive closure of the relation *refers to*.
3. Neg_P denotes the set of predicate symbols in P which occur in a negative literal in the body of a clause in P.
4. Neg_P^* denotes the set of all predicate symbols in P on which the predicate symbols in Neg_P depend.
5. P^- denotes the set of clauses in P whose head contains a predicate symbol from Neg_P^*.

Definition 2 A program P is called Φ^*-*accessible* if and only if there exists a level mapping l for P and a model I for P which is a supported model of P^-, such that the following condition holds. For each clause $A \leftarrow L_1, \ldots, L_n$ in ground(P), we either have $I \models L_1 \wedge \cdots \wedge L_n$ and $l(A) > l(L_i)$ for all $i = 1, \ldots, n$ or there exists $i \in \{1, \ldots, n\}$ such that $I \not\models L_i$ and $l(A) > l(L_i)$.

The Φ^*-accessible programs are a common generalization of acyclic, locally hierarchical [3] and acceptable [2] programs. In [6], the authors gave a unified treatment of these classes of programs by means of operators in various three-valued logics.

Definition 3 Let X be a set and let Γ be a partially ordered set with least element 0. We call

(X, d, Γ) (or simply (X, d)) a *generalized ultrametric space* if $d : X \times X \to \Gamma$ is a function such that for all $x, y, z \in X$ and all $\gamma \in \Gamma$ we have:

(Ui) $d(x, y) = 0$ implies $x = y$.
(Uii) $d(x, x) = 0$.
(Uiii) $d(x, y) = d(y, x)$.
(Uiv) Whenever $d(x, y) \leq \gamma$ and $d(y, z) \leq \gamma$, we have $d(x, z) \leq \gamma$.

Generalized ultrametrics have been studied in the context of logic programming semantics in [7, 15, 16]. If d satisfies conditions (Ui), (Uiii) and (Uiv) only, we call (X, d) a *dislocated generalized ultrametric space* or simply a *d-ultrametric space*.

A Generalized Priess-Crampe & Ribenboim Fixed-Point Theorem

Definition 4 Let (X, d, Γ) be a d-ultrametric space. For $0 \neq \gamma \in \Gamma$ and $x \in X$, the set $B_\gamma(x) := \{y \in X \mid d(x, y) \leq \gamma\}$ is called a (γ-)ball in X with centre x. A d-ultrametric space is called *spherically complete* if, for any chain $(\mathcal{C}, \subseteq)$ of non-empty balls in X, we have $\bigcap \mathcal{C} \neq \emptyset$. A function $f : X \to X$ is called

(1) *non-expanding* if $d(f(x), f(y)) \leq d(x, y)$ for all $x, y \in X$,
(2) *strictly contracting on orbits* if $d(f^2(x), f(x)) < d(f(x), x)$ for every $x \in X$ with $x \neq f(x)$, and
(3) *strictly contracting* if $d(f(x), f(y)) < d(x, y)$ for all $x, y \in X$ with $x \neq y$.

We will need the following observations, which are well-known for ordinary ultrametric spaces, see [14].

Lemma 5 Let (X, d, Γ) be a d-ultrametric space. For $\alpha, \beta \in \Gamma$ and $x, y \in X$ the following statements hold.

(1) If $\alpha \leq \beta$ and $B_\alpha(x) \cap B_\beta(y) \neq \emptyset$, then $B_\alpha(x) \subseteq B_\beta(y)$.
(2) If $B_\alpha(x) \cap B_\alpha(y) \neq \emptyset$, then $B_\alpha(x) = B_\alpha(y)$.
(3) $B_{d(x,y)}(x) = B_{d(x,y)}(y)$.

Proof: Let $a \in B_\alpha(x)$ and $b \in B_\alpha(x) \cap B_\beta(y)$. Then $d(a, x) \leq \alpha$ and $d(b, x) \leq \alpha$, hence $d(a, b) \leq \alpha \leq \beta$. Since $d(b, y) \leq \beta$, we have $d(a, y) \leq \beta$, hence $a \in B_\beta(y)$, which proves the first statement. The second follows by symmetry and the third by replacing α by $d(x, y)$. ■

The following theorem gives a partial unification of a theorem of Matthews [12, Theorem 5, Page 20] and the Priess-Crampe & Ribenboim theorem. The proof of the latter theorem given in [14] in fact carries over directly to our more general setting of d-ultrametrics.

Theorem 6 Let (X, d, Γ) be a spherically complete d-ultrametric space and let $f : X \to X$ be non-expanding and strictly contracting on orbits. Then f has a fixed point. If f is strictly contracting on X, then the fixed point is unique.

Proof: Assume that f has no fixed point. Then for all $x \in X, d(x, f(x)) \neq 0$. We define the set $\mathcal{B}$ by $\mathcal{B} = \{B_{d(x,f(x))}(x) \mid x \in X\}$, and note that each ball in this set is non-empty. We also note that because $d(x, y)) \leq d(x, y) = d(y, x)$, and using (Uiv), we have $d(x, x) \leq d(x, y)$ for all $x, y \in X$, and it follows easily from this that $B_{d(x,f(x))}(x) = B_{d(x,f(x))}(f(x))$ by Lemma 5. Now let $\mathcal{C}$ be a maximal chain in $\mathcal{B}$. Since X is spherically complete, there exists $z \in \bigcap \mathcal{C}$. We show that $B_{d(z,f(z))}(z) \subseteq B_{d(x,f(x))}$ for all $x \in X$ and hence, by maximality, that $B_{d(z,f(z))}(z)$ is the smallest ball in the chain. Let $B_{d(x,f(x))}(x) \in \mathcal{C}$. Since $z \in B_{d(x,f(x))}(x)$, and noting our earlier observation that $B_{d(x,f(x))}(x) = B_{d(x,f(x))}(f(x))$ for all x, we get $d(z, x) \leq d(x, f(x))$ and $d(z, f(x)) \leq d(x, f(x))$. By non-expansiveness of f, we get $d(f(z), f(x)) \leq d(z, x) \leq d(x, f(x))$. It follows by (Uiv) that $d(z, f(z)) \leq d(x, f(x))$ and therefore that $B_{d(z,f(z))}(z) \subseteq B_{d(x,f(x))}(x)$ by Lemma 5 for all $x \in X$, since x was chosen arbitrarily. Now, since f is strictly contracting on orbits, $d(f(z), f^2(z)) < d(z, f(z))$, and therefore $z \notin B_{d(f(z),f^2(z))}(f(z)) \subset B_{d(z,f(z))}(f(z))$. By Lemma 5, this is equivalent to $B_{d(f(z),f^2(z))}(f(z)) \subset B_{d(z,f(z))}(z)$, which is a contradiction to the maximality of $\mathcal{C}$. So f has a fixed point.

Now let f be strictly contracting on X and assume that x, y are two distinct fixed points of

f. Then we get $d(x,y) = d(f(x), f(y)) < d(x,y)$ which is impossible. So the fixed point of f is unique in this case. ■

We note that if d is a d-ultrametric, we can generate an associated generalized ultrametric d' in the usual sense by defining $d'(x,y) = d(x,y)$ whenever $x \neq y$ and setting $d'(x,x) = 0$ for all x. Doing this, however, does not simplify our main application, which is below, and one then has to check that spherical completeness is preserved in generalizing the theorem of Priess-Crampe & Ribenboim. Since distance functions d such that $d(x,x)$ is not necessarily equal to 0 do arise naturally in computing, and we consider one next, we prefer to stay with the d-ultrametric and not pass to an ultrametric.

An Application of Theorem 6: Φ^*-accessible programs

In the following, P is a Φ^*-accessible program which satisfies the defining conditions with respect to a model I and a level mapping $l : B_P \to \gamma$. We let Γ denote the set $\{2^{-\alpha} \mid \alpha \leq \gamma\}$ ordered by $2^{-\alpha} < 2^{-\beta}$ iff $\beta < \alpha$, and denote $2^{-\gamma}$ by 0.

For $J, K \in I_P$, we now define $d(K,K) = 0$, and $d(J,K) = 2^{-\alpha}$, where J and K differ on some atom $A \in B_P$ of level α, but agree on all ground atoms of lower level. As was shown in [16], (I_P, d) is a spherically complete generalized ultrametric space. For $K \in I_P$, we denote by K' the set K restricted to the predicate symbols in Neg_P^*. By analogy with [5], we now define for all $J, K \in I_P$: $d_1(J,K) = d(J', K')$ and $d_2(J,K) = d(J \setminus J', K \setminus K')$. Finally, define a function $f : I_P \to \Gamma$ by $f(K) = 0$ if $K \setminus K' \subseteq I$ and otherwise $f(K) = 2^{-\alpha}$, where α is the smallest ordinal such that there is an atom $A \in K \setminus K'$ with $l(A) = \alpha$ and $A \notin I$. Finally, we define $\varrho(J,K) = \max\{d_1(J,I), d_1(K,I), d_2(J,K), f(J), f(K)\}$ for all $J, K \in I_P$.

Proposition 7 (I_P, ϱ) is a spherically complete d-ultrametric space.

Proof: (Ui), (Uiii) and (Uiv) we leave to the reader. For spherical completeness, let $(\mathcal{B}_\alpha)$ be a (decreasing) chain of balls in X with centres I_α. Let K be the set of all atoms which are *eventually* in I_α, that is, the set of all $A \in B_P$ such that there exists some β with $A \in I_\alpha$ for all $\alpha \geq \beta$. We show that for each ball $B_{2^{-\alpha}}(I_\alpha)$ in the chain we have $d(I_\alpha, I) \leq 2^{-\alpha}$, which suffices to show that K is in the intersection of the chain. Indeed, it is easy to see by the definition of ϱ that all I_β with $\beta > \alpha$ agree on all atoms of level less than α. Hence, by definition of K we obtain that K and I_α agree on all atoms of level less than α as required. ■

The next proposition is analogous to [5, Proposition 7.1].

Proposition 8 Let P be Φ^*-accessible with respect to a level mapping l and a model I. Then for all $J, K \in I_P$ with $J \neq K$ we have $\varrho(T_P(J), T_P(K)) < \varrho(J,K)$. In particular we have the following:

(i) $d_1(T_P(J), I) < d_1(J, I)$.
(ii) $f(T_P(J)), f(T_P(K)) < \varrho(J,K)$.
(iii) $d_2(T_P(J), T_P(K)) < \varrho(J,K)$.

Proof: It suffices to prove properties (i), (ii) and (iii). For convenience, we identify Neg_P^* with the subset of B_P containing predicate symbols from Neg_P^*.

(i) First note that $d_1(T_P(J), I) = d_1(T_{P^-}(J), I)$ since d_1 only depends on the predicate symbols in Neg_P^*. Let $d(J,I) = 2^{-\alpha}$. We show that $d(T_{P^-}(J), I) \leq 2^{-(\alpha+1)}$. We know that J' and I' agree on all ground atoms of level less than α and differ on an atom of level α. It suffices to show now that $T_{P^-}(J)'$ and I' agree on all ground atoms of level less than or equal to α.

Let A be a ground atom in Neg_P^* with $l(A) \leq \alpha$ and suppose that $T_{P^-}(J)$ and I differ on A. Assume first that $A \in T_{P^-}(J)$ and $A \notin I$. Then there must be a ground instance $A \leftarrow L_1, \ldots, L_m$ of a clause in P^- such that $J \models L_1 \wedge \cdots \wedge L_m$. Since I is a fixed point of T_{P^-}, and using Definition 2, there must also be a k such that $I \not\models L_k$ and $l(L_k) < \alpha$. Note that the predicate symbol in L_k is contained in Neg_P^*. So we obtain $I \not\models L_k$, $J \models L_k$ and $l(L_k) < \alpha$ which is a contradiction to the assumption that J and I agree on all atoms in Neg_P^* of level less than α. Now assume that $A \in I$ and $A \notin T_{P^-}(J)$. It follows that there is a ground instance $A \leftarrow L_1, \dots, L_m$ of a clause in P^- such that $I \models L_1 \wedge \cdots \wedge L_m$ and $l(A) > l(L_1), \dots, l(L_m)$ by Definition 2. But then $J \models L_1 \wedge \cdots \wedge L_m$ since J and I agree on all atoms of level less than α and consequently $A \in T_{P^-}(J)$. This contradiction establishes (i).

(ii) It suffices to show this for K. Assume $\varrho(J,K) = 2^{-\alpha}$. We show that $f(T_P(K)) \leq 2^{-(\alpha+1)}$, for which in turn we have to show that for each $A \in T_P(K)$ not in Neg_P^* with $l(A) \leq \alpha$ we have $A \in I$. Assume that $A \notin I$ for such an A. Since $A \in T_P(K)$, there is a ground instance $A \leftarrow L_1, \dots, L_m$ of a clause in P with $K \models L_1 \wedge \cdots \wedge L_m$. Since $A \notin I$, there must also be a k with $I \not\models L_k$ and $l(A) > l(L_k)$ by Definition 2. If the predicate symbol of L_k belongs to Neg_P^* then, since K and I agree on all atoms in Neg_P^* of level less than α, we obtain $K \not\models L_k$ which contradicts $K \models L_1 \wedge \cdots \wedge L_m$. If the predicate symbol in L_k does not belong to Neg_P^*, then L_k is an atom and since $f(K) \leq 2^{-\alpha}$ we obtain $I \models L_k$, which is again a contradiction.

(iii) Let $\varrho(J,K) = 2^{-\alpha}$, let A be not in Neg_P^* with $l(A) \leq \alpha$ and $A \in T_P(J)$. By symmetry, it suffices to show that $A \in T_P(K)$. Since $A \in T_P(J)$, we must have a ground instance $A \leftarrow L_1, \dots, L_m$ of a clause in P with $J \models L_1 \wedge \cdots \wedge L_m$. If $I \models L_1 \wedge \cdots \wedge L_m$, then $l(L_k) < l(A) \leq \alpha$ for all k, and since J and K agree on all atoms of level less than α we obtain $K \models L_1 \wedge \cdots \wedge L_m$, and hence $A \in T_P(K)$. If there is some L_k such that $I \not\models L_k$, then without loss of generality $l(L_k) < l(A) \leq \alpha$ by Definition 2. Now, if the predicate symbol of L_k belongs to Neg_P^* then, since $d_1(J,I) \leq 2^{-\alpha}$, we obtain from $J \models L_k$ that $I \models L_k$ which is a contradiction. Also, if the predicate symbol of L_k does not belong to Neg_P^*, then L_k is an atom and since $f(J) \leq 2^{-\alpha}$, we obtain $I \models L_k$, again a contradiction. This establishes (iii). ∎

We are now in a position to prove our main result.

Theorem 9 Let P be a Φ^*-accessible program. Then P has a unique supported model.

Proof: By Proposition 8, T_P is strictly contracting with respect to d_3, which in turn is a spherically complete d-ultrametric by Proposition 7. So by Theorem 6, the operator T_P must have a unique fixed point which yields a unique supported model for P. ∎

Conclusion

This work is part of an ongoing programme of research being undertaken by the authors in investigating the extent to which the methods of domain theory and denotational semantics can be applied to logic programming, nonmonotonic reasoning and artificial intelligence. The focus of much of this work is on the fixed points of various operators which are associated with programs written in these paradigms, since the former provide one with a semantics (the fixed-point semantics) for the latter. Such methods depend heavily on ideas and techniques drawn from topology as illustrated in this paper, and this line of research is being further pursued in [8, 9].

References

[1] K.R. Apt, H.A. Blair and A. Walker, “Towards a Theory of Declarative Knowledge”, in: J. Minker (Ed.), Foundations of Deductive Databases and Logic Programming, Morgan Kaufmann, Los Altos, CA, 1988, pp. 89–148.

[2] K.R. Apt and D. Pedreschi, “Reasoning about Termination of Pure Prolog Programs”, Information and Computation **106** (1993), pp. 109–157.

[3] L. Cavedon, “Acyclic Programs and the Completeness of SLDNF-Resolution”, Theoretical Computer Science **86** (1991), pp. 81–92.

[4] K.L. Clark, “Negation as Failure”, in: H. Gallaire and J. Minker (Eds.), Logic and Data Bases, Plenum Press, New York, 1978, pp. 293–322.

[5] M. Fitting, “Metric Methods: Three Examples and a Theorem”, J. Logic Programming **21** (3) (1994), pp. 113–127.

[6] P. Hitzler and A.K. Seda, “Characterizations of Classes of Programs by Three-Valued Operators”, in: Gelfond, M., Leone,

N. and Pfeifer, G. (Eds.), Logic Programming and Nonmonotonic Reasoning, Proceedings of the 5th International Conference on Logic Programming and Non-Monotonic Reasoning (LPNMR'99), El Paso, Texas, USA, December 1999. Lecture Notes in Artificial Intelligence, Vol. 1730, Springer, Berlin, 1999, pp. 357–371.

[7] P. Hitzler and A.K. Seda, "The Fixed-Point Theorems of Priess-Crampe and Ribenboim in Logic Programming", Proceedings of the International Conference and Workshop on Valuation Theory, University of Saskatchewan in Saskatoon, Canada, July 1999. Fields Institute Communications Series, American Mathematical Society, Providence, R.I., 17 pages, to appear.

[8] P. Hitzler and A.K. Seda, "Some Issues Concerning Fixed Points in Computational Logic: Quasi-Metrics, Multivalued Mappings and the Knaster-Tarski Theorem", Proceedings of the 14th Summer Conference on Topology and its Applications: Special Session on Topology in Computer Science, New York, August 1999. Topology Proceedings, pp. 1-15, to appear.

[9] P. Hitzler and A.K. Seda, "Dislocated Topologies", Proceedings of the Slovakian Conference in Applied Mathematics 2000, Bratislava, April 2000, Journal of Electrical Engineering, Slovak Academy of Sciences, pp. 1-8, to appear.

[10] P. Hitzler and A.K. Seda, "Classes of Logic Programs which Possess Unique Supported Models", Preprint, Department of Mathematics, University College Cork, Cork, Ireland, 2000, pp. 1-15.

[11] J.W. Lloyd, Foundations of Logic Programming, Berlin: Springer, 1988.

[12] S.G. Matthews, "Metric Domains for Completeness", Ph.D. Thesis, 1985. Research Report 76, Department of Computer Science, University of Warwick, UK, April 1986.

[13] S.G. Matthews, "Partial Metric Topology", Proceedings of the Eighth Summer Conference on General Topology and its Applications, Annals of the New York Academy of Sciences, Vol. 728, New York, 1994, pp. 183-197.

[14] S. Priess-Crampe and P. Ribenboim, "Fixed Points, Combs and Generalized Power Series", Abh. Math. Sem. Univ. Hamburg **63** (1993), pp. 227–244.

[15] S. Priess-Crampe and P. Ribenboim, "Ultrametric Spaces and Logic Programming", J. Logic Programming **42** (2000), pp. 59–70.

[16] A.K. Seda and P. Hitzler, "Topology and Iterates in Computational Logic", Proc. 12th Summer Conference on General Topology and its Applications, Ontario, 1997. Topology Proceedings Vol. 22, Summer 1997, pp. 427-469.