

The Category-Theoretic Solution of Recursive Metric-Space Equations

Lars Birkedal* Kristian Støvring†

Jacob Thamsborg‡

IT University of Copenhagen§

July 21, 2010

Abstract

It is well known that one can use an adaptation of the inverse-limit construction to solve recursive equations in the category of complete ultrametric spaces. We show that this construction generalizes to a large class of categories with metric-space structure on each set of morphisms: the exact nature of the objects is less important. In particular, the construction immediately applies to categories where the objects are ultrametric spaces with ‘extra structure’, and where the morphisms preserve this extra structure. The generalization is inspired by classical domain-theoretic work by Smyth and Plotkin.

For many of the categories we consider, there is a natural subcategory in which each set of morphisms is required to be a compact metric space. Our setting allows for a proof that such a subcategory always inherits solutions of recursive equations from the full category.

As another application, we present a construction that relates solutions of generalized domain equations in the sense of Smyth and Plotkin to solutions of equations in our class of categories.

Our primary motivation for solving generalized recursive metric-space equations comes from recent and ongoing work on Kripke-style models in which the sets of worlds must be recursively defined. We show a series of examples motivated by this line of work.

Keywords: Metric space, fixed point, recursive equation.

1 Introduction

Smyth and Plotkin [23] showed that in the classical inverse-limit construction of solutions to recursive domain equations, what matters is not that the *objects* of the category under consideration are domains, but that the sets of *morphisms* between objects are domains. In this article we show that, in the case of ultrametric spaces, the standard construction of solutions to recursive metric-space equations [6, 11] can be similarly generalized to a large class of categories with metric-space structure on each set of morphisms. The generalization in particular allows one to solve recursive

*Email: birkedal@itu.dk.

†Corresponding author. Email: kss@itu.dk. Phone: (+45)72185047. Fax: (+45)72185001.

‡Email: thamsborg@itu.dk

§Rued Langgaards Vej 7, 2300 Copenhagen S, Denmark.

equations in categories where the objects are ultrametric spaces with some form of additional structure, and where the morphisms preserve this additional structure. Some applications from recent and ongoing work in semantics are shown in Section 7.

For many of the categories we consider, there is a natural variant, indeed a subcategory, in which each set of morphisms is required to be a compact metric space [3, 10]. Our setting allows for a general proof that such a subcategory inherits solutions of recursive equations from the full category. Otherwise put, the problem of solving recursive equations in such a ‘locally compact’ subcategory is, in a certain sense, reduced to the similar problem for the full category. The fact that one can solve recursive equations in a category of compact ultrametric spaces [10] arises as a particular instance. (For various applications of compact metric spaces in semantics, see the references in the introduction to van Breugel and Warmerdam [10].)

As another application, we present a construction that relates solutions of generalized domain equations in the sense of Smyth and Plotkin to solutions of equations in our class of categories. This construction generalizes and improves an earlier one due to Baier and Majster-Cederbaum [7].

The key to achieving the right level of generality in the results lies in inspiration from enriched category theory. We shall not refer to general enriched category theory below, but rather present the necessary definitions in terms of metric spaces. The basic idea is, however, that given a cartesian category V (or more generally, a monoidal category), one considers so-called V -categories, in which the ‘hom-sets’ are in fact objects of V instead of sets, and where the ‘composition functions’ are morphisms in V .

Other related work. The idea of considering categories with metric spaces as hom-sets has been used in earlier work [10, 20]. Rutten and Turi [20] show existence and uniqueness of fixed points in a particular category of (not necessarily ultrametric) spaces, but with a proof where parts are more general: some aspects of our Lemma 3.2 are covered. In other work, van Breugel and Warmerdam [10] show uniqueness for a more general notion of categories than ours, again not requiring ultrametricity. Neither of these articles contain a theorem about existence of fixed points for a general class of ‘metric-enriched’ categories (as in our Theorem 3.1), nor a general theorem about fixed points in locally compact subcategories (Theorem 4.1).

Alessi et al. [4] consider solutions to *non-functorial* recursive equations in certain categories of metric spaces, i.e., recursive equations whose solutions cannot necessarily be described as fixed-points of functors. In contrast, we only consider *functorial* recursive equations in this article.

Wagner [27] gives a comprehensive account of a generalized inverse limit construction that in particular works for categories of metric spaces and categories of domains. Our generalization is in a different direction, namely to categories where the hom-sets are metric spaces. We do not know whether there is a common generalization of our work and Wagner’s work. In this article we do not aim for maximal generality, but rather for a level of generality that seems right for applications in the style of those in Section 7.

A more detailed discussion of the level of generality of our results, and of their relation to results in the literature, can be found in Section 9.

2 Ultrametric spaces

We first recall some basic definitions and properties about metric spaces [22].

A metric space (X, d) is *1-bounded* if $d(x, y) \leq 1$ for all x and y in X . We shall only work with 1-bounded metric spaces. One advantage of doing so is that one

can define coproducts and general products of such spaces; alternatively, one could have allowed infinite distances.

An *ultrametric space* is a metric space (X, d) that satisfies the ‘ultrametric inequality,’

$$d(x, z) \leq \max(d(x, y), d(y, z)),$$

and not just the weaker triangle inequality (where one has $+$ instead of \max on the right-hand side). It might be helpful to think of the function d of an ultrametric space (X, d) not as a measure of (euclidean) distance between elements, but rather as a measure of the degree of similarity between elements.

A function $f : X_1 \rightarrow X_2$ from a metric space (X_1, d_1) to a metric space (X_2, d_2) is *non-expansive* if $d_2(f(x), f(y)) \leq d_1(x, y)$ for all x and y in X_1 . Stronger, such a function f is *contractive* if there exists $c < 1$ such that $d_2(f(x), f(y)) \leq c \cdot d_1(x, y)$ for all x and y in X_1 . Notice that a non-expansive function is (uniformly) continuous in the metric-space sense and hence preserves limits of convergent sequences.

A metric space is *complete* if it is Cauchy-complete in the usual sense, i.e., if every Cauchy sequence in the metric space has a limit. By Banach’s fixed-point theorem, every contractive function from a non-empty, complete metric space to itself has a unique fixed point.

In the following we only consider complete, 1-bounded ultrametric spaces. As a canonical example of such a metric space, consider the set \mathbb{N}^ω of infinite sequences of natural numbers, with distance function d given by:

$$d(x, y) = \begin{cases} 2^{-\max\{n \in \omega \mid \forall m \leq n. x(m) = y(m)\}} & \text{if } x \neq y \\ 0 & \text{if } x = y. \end{cases}$$

To avoid confusion, call the elements of \mathbb{N}^ω *strings* instead of sequences. Here the ultrametric inequality simply states that if x and y agree on the first n ‘characters’ and y and z also agree on the first n characters, then x and z agree on the first n characters. A Cauchy sequence in \mathbb{N}^ω is a sequence of strings $(x_n)_{n \in \omega}$ in which the individual characters ‘stabilize’: for every m there exists $N \in \omega$ such that $x_{n_1}(m) = x_{n_2}(m)$ for all $n_1, n_2 \geq N$.

Let \mathbf{CBUIt} be the category with complete, 1-bounded ultrametric spaces as objects and non-expansive functions as morphisms [6]. This category is cartesian closed [22]; here one needs the ultrametric inequality. The terminal object is the one-point metric space. Binary products are defined in the natural way: $(X_1, d_1) \times (X_2, d_2) = (X_1 \times X_2, d_{X_1 \times X_2})$ where

$$d_{X_1 \times X_2}((x_1, x_2), (y_1, y_2)) = \max(d_1(x_1, y_1), d_2(x_2, y_2)).$$

The exponential $(X_1, d_1) \rightarrow (X_2, d_2)$, sometimes written $(X_2, d_2)^{(X_1, d_1)}$, has the set of non-expansive functions from (X_1, d_1) to (X_2, d_2) as the underlying set, and the ‘sup’-metric $d_{X_1 \rightarrow X_2}$ as distance function: $d_{X_1 \rightarrow X_2}(f, g) = \sup\{d_2(f(x), g(x)) \mid x \in X_1\}$. For both products and exponentials, limits are pointwise. It follows from the cartesian closed structure that the function $(X_3, d_3)^{(X_2, d_2)} \times (X_2, d_2)^{(X_1, d_1)} \rightarrow (X_3, d_3)^{(X_1, d_1)}$ given by composition is non-expansive; this fact is needed in several places below.

Moreover, the category \mathbf{CBUIt} is complete [19]: general products are defined in the same way as binary ones, except that the distance function on an infinite product space is in general given by a supremum instead of a maximum. An equalizer of two parallel arrows $f, g : X \rightarrow Y$ is given by the subset $\{x \in X \mid f(x) = g(x)\}$ of X , with the metric inherited from X .

\mathbf{CBUIt} is also cocomplete. The coproduct of a family $(X_j, d_j)_{j \in J}$ of \mathbf{CBUIt} -objects is $(\coprod_{j \in J} X_j, d)$ where $\coprod_{j \in J} X_j$ is the disjoint union (coproduct in \mathbf{Set}) of

the underlying sets X_j , and where the distance function d is given by

$$d(x, y) = \begin{cases} d_j(x, y), & \text{if } x \in X_j \text{ and } y \in X_j \text{ for some } j \in J, \\ 1, & \text{otherwise.} \end{cases}$$

Coequalizers are more complicated to describe, and we shall not need them in this article.

It is a trivial fact, but for our purposes a rather annoying one, that Banach's fixed-point theorem only holds for non-empty metric spaces. To avoid tedious special cases below, we shall therefore not work with the category \mathbf{CBUlt} , but rather with the full subcategory $\mathbf{CBUlt}_{\text{ne}}$ of *non-empty*, complete, 1-bounded ultrametric spaces. This category is also cartesian closed: since it is a full subcategory of \mathbf{CBUlt} , it suffices to verify that \mathbf{CBUlt} -products of non-empty metric spaces are non-empty, and similarly for exponentials. The category $\mathbf{CBUlt}_{\text{ne}}$ is not complete, and in fact it does not even have all limits of ω^{op} -chains. We return to that point in Section 5.

In some settings it is useful to work with *compact* metric spaces [3, 10]. Recall that a metric space is compact in the usual topological sense if and only if it is both complete and *totally bounded* [22]: for every $\epsilon > 0$, there exist finitely many points x_1, \dots, x_n in the space such that the open balls with centers x_i and radius ϵ cover the space. As a canonical example of a compact, 1-bounded ultrametric space, consider the set $\{0, 1\}^\omega$ of infinite sequences of zeros or ones, with distance function given as in the example with sequences of natural numbers above. (Any finite set other than $\{0, 1\}$ would also work.)

Let \mathbf{KBUlt} be the full subcategory of \mathbf{CBUlt} consisting of compact, 1-bounded ultrametric spaces, and let $\mathbf{KBUlt}_{\text{ne}}$ be the full subcategory of *non-empty* such spaces. Both of these categories are cartesian closed [22] and have finite coproducts. \mathbf{KBUlt} has all finite limits, but neither \mathbf{KBUlt} nor $\mathbf{KBUlt}_{\text{ne}}$ is complete. We return to that point in Section 4.

2.1 M -categories

Recall from the introduction that the basic idea of this article is to generalize a theorem about a particular category of metric spaces, here $\mathbf{CBUlt}_{\text{ne}}$, to a theorem about all ' $\mathbf{CBUlt}_{\text{ne}}$ -categories' where the hom-sets are in fact appropriate metric spaces. In analogy with the O -categories of Smyth and Plotkin (O for 'order' or 'ordered') we call such categories M -categories.

Definition 2.1. An M -category is a category \mathcal{C} where each hom-set $\mathcal{C}(A, B)$ is equipped with a distance function turning it into a non-empty, complete, 1-bounded ultrametric space, and where each composition function $\circ : \mathcal{C}(B, C) \times \mathcal{C}(A, B) \rightarrow \mathcal{C}(A, C)$ is non-expansive with respect to these metrics. (Here the domain of such a composition function is given the product metric.)

In other words, an M -category is a category where each hom-set is equipped with a metric which turns it into an object in $\mathbf{CBUlt}_{\text{ne}}$; furthermore, each composition function must be a morphism in $\mathbf{CBUlt}_{\text{ne}}$.

A simple example of an M -category is $\mathbf{CBUlt}_{\text{ne}}$ itself. The distance function on each hom-set $\mathbf{CBUlt}_{\text{ne}}(A, B)$ is defined as for the exponential B^A in $\mathbf{CBUlt}_{\text{ne}}$, i.e., $d(f, g) = \sup\{d_B(f(x), g(x)) \mid x \in A\}$. The fact that the composition functions are non-expansive, as observed in Section 2, depends on the ultrametric inequality. Since $\mathbf{CBUlt}_{\text{ne}}$ is itself an M -category the results below can be used to solve standard recursive equations over ultrametric spaces.

Let \mathcal{C} be an M -category. A functor $F : \mathcal{C}^{\text{op}} \times \mathcal{C} \rightarrow \mathcal{C}$ is *locally non-expansive* if $d(F(f, g), F(f', g')) \leq \max(d(f, f'), d(g, g'))$ for all f, f', g , and g' with appropriate

domains and codomains. In other words, such an F is locally non-expansive if each component

$$F_{A,A',B,B'} : \mathcal{C}(A', A) \times \mathcal{C}(B, B') \rightarrow \mathcal{C}(F(A, B), F(A', B'))$$

is a morphism in $\mathbf{CBUlt}_{\text{ne}}$. Stronger, F is *locally contractive* if there exists some $c < 1$ such that $d(F(f, g), F(f', g')) \leq c \cdot \max(d(f, f'), d(g, g'))$ for all f, f', g , and g' . Notice that c is global in the sense that it is a common ‘contractiveness factor’ for all components of the functor: each component $F_{A,A',B,B'}$ is contractive with factor c .

In the particular categories we consider in the examples in Section 7, many ‘natural’ functors such as those given by binary products or coproducts are only locally non-expansive, not locally contractive. On each of these categories \mathcal{C} there is, however, an appropriate functor $\frac{1}{2} : \mathcal{C} \rightarrow \mathcal{C}$ which multiplies all distances in hom-sets by the factor $1/2$. Composing a locally non-expansive functor with $\frac{1}{2}$ then yields a locally contractive functor.

3 Solving recursive equations

Let \mathcal{C} be an M -category. We consider mixed-variance functors $F : \mathcal{C}^{\text{op}} \times \mathcal{C} \rightarrow \mathcal{C}$ on \mathcal{C} and recursive equations of the form

$$X \cong F(X, X).$$

In other words, given such an F we seek a fixed point of F up to isomorphism.

Covariant endofunctors on \mathcal{C} are a special case of mixed-variance functors. It would in some sense suffice to study covariant functors: if \mathcal{C} is an M -category, then so are \mathcal{C}^{op} (with the same metric on each hom-set as in \mathcal{C}) and $\mathcal{C}^{\text{op}} \times \mathcal{C}$ (with the product metric on each hom-set), and it is well-known how to construct a ‘symmetric’ endofunctor on $\mathcal{C}^{\text{op}} \times \mathcal{C}$ from a functor such as F above. We explicitly study mixed-variance functors since the proof of the existence theorem below would in any case involve an M -category of the form $\mathcal{C}^{\text{op}} \times \mathcal{C}$. As a benefit we directly obtain theorems of the form useful in applications. For example, for the existence theorem we are interested in completeness conditions on \mathcal{C} , not on $\mathcal{C}^{\text{op}} \times \mathcal{C}$.

3.1 Uniqueness of solutions

The results below depend on the assumption that the given functor F is locally contractive. One easy consequence of this assumption is that, unlike in the domain-theoretic setting [23], there is at most one fixed point of F up to isomorphism.

Theorem 3.1. Let $F : \mathcal{C}^{\text{op}} \times \mathcal{C} \rightarrow \mathcal{C}$ be a locally contractive functor on an M -category \mathcal{C} , and assume that $i : F(A, A) \rightarrow A$ is an isomorphism. Then the pair (i, i^{-1}) is a *bifree algebra* for F in the following sense: for all objects B of \mathcal{C} and all morphisms $f : F(B, B) \rightarrow B$ and $g : B \rightarrow F(B, B)$, there exists a unique pair of morphisms $(k : B \rightarrow A, h : A \rightarrow B)$ such that $h \circ i = f \circ F(k, h)$ and $i^{-1} \circ k = F(h, k) \circ g$:

$$\begin{array}{ccc}
 F(A, A) & \begin{array}{c} \xrightarrow{F(k, h)} \\ \xleftarrow{F(h, k)} \end{array} & F(B, B) \\
 \begin{array}{c} \uparrow \\ i^{-1} \end{array} \downarrow i & & \begin{array}{c} \uparrow \\ g \end{array} \downarrow f \\
 A & \begin{array}{c} \text{---} \xrightarrow{h} \text{---} \\ \text{---} \xleftarrow{k} \text{---} \end{array} & B
 \end{array}$$

In particular, A is the unique fixed point of F up to isomorphism.

Proof. First we observe that there is an obvious way to define a category of ‘bialgebras’ for F such that a bifree algebra, as defined above, is an initial object in this category. It follows that any two bifree algebras are isomorphic as bialgebras, hence that their underlying \mathcal{C} -objects are isomorphic. So once we have shown that every fixed point of F is a bifree algebra, it follows that there is at most one fixed point of F up to isomorphism.

Let now $i : F(A, A) \rightarrow A$ be an isomorphism and assume that F is locally contractive with factor $c < 1$; we show that (i, i^{-1}) is a bifree algebra. Let $f : F(B, B) \rightarrow B$ and $g : B \rightarrow F(B, B)$ be given. Recall that hom-sets in \mathcal{C} are equipped with metrics that turn them into objects of CBUlt_{ne} , and let X be the CBUlt_{ne} -object $\mathcal{C}(B, A) \times \mathcal{C}(A, B)$. We obtain the desired pair of morphisms $(k : B \rightarrow A, h : A \rightarrow B)$ as the unique fixed point of the following contractive operator on X :

$$D(k, h) = (i \circ F(h, k) \circ g, f \circ F(k, h) \circ i^{-1}).$$

First we verify that this operator is indeed contractive. Given (k_1, h_1) and (k_2, h_2) in X ,

$$\begin{aligned} d(D(k_1, h_1), D(k_2, h_2)) &= \max(d(i \circ F(h_1, k_1) \circ g, i \circ F(h_2, k_2) \circ g), \\ &\quad d(f \circ F(k_1, h_1) \circ i^{-1}, f \circ F(k_2, h_2) \circ i^{-1})), \end{aligned}$$

by the definition of the product metric. But the composition functions of an M -category are required to be non-expansive: therefore,

$$\begin{aligned} d(i \circ F(h_1, k_1) \circ g, i \circ F(h_2, k_2) \circ g) &\leq \max(d(i, i), d(F(h_1, k_1), F(h_2, k_2)), d(g, g)) \\ &= d(F(h_1, k_1), F(h_2, k_2)) \\ &\leq c \cdot \max(d(h_1, h_2), d(k_1, k_2)) \\ &= c \cdot d((k_1, h_1), (k_2, h_2)), \end{aligned}$$

and similarly,

$$\begin{aligned} d(f \circ F(k_1, h_1) \circ i^{-1}, f \circ F(k_2, h_2) \circ i^{-1}) &\leq d(F(k_1, h_1), F(k_2, h_2)) \\ &\leq c \cdot \max(d(h_1, h_2), d(k_1, k_2)) \\ &= c \cdot d((k_1, h_1), (k_2, h_2)). \end{aligned}$$

Therefore, $d(D(k_1, h_1), D(k_2, h_2)) \leq c \cdot d((k_1, h_1), (k_2, h_2))$, and D is locally contractive with factor c .

Since hom-sets of \mathcal{C} are *non-empty* complete metric spaces, the operator D has a unique fixed point by Banach’s theorem. It only remains to show that a pair of morphisms $(k : B \rightarrow A, h : A \rightarrow B)$ is a fixed point of D if and only if it makes the diagram in the statement of the theorem commute. But this is easy since i is an isomorphism: $k = i \circ F(h, k) \circ g$ holds if and only if $i^{-1} \circ k = F(h, k) \circ g$ holds, and similarly, $h = f \circ F(k, h) \circ i^{-1}$ holds if and only if $h \circ i = f \circ F(k, h)$ holds. We conclude that (i, i^{-1}) is a bifree algebra for F . \square

In particular, if F is covariant and $i : FA \rightarrow A$ is an isomorphism, then i is an initial F -algebra and i^{-1} is a final F -coalgebra. As an example, consider the M -category CBUlt_{ne} and take F to be the covariant functor $\frac{1}{2} : \text{CBUlt}_{\text{ne}} \rightarrow \text{CBUlt}_{\text{ne}}$ which given a metric space yields the same metric space but with all distances multiplied by $1/2$, and which is the identity on morphisms. Evidently, the one-point metric space is a fixed-point of F . By the theorem above it is also an initial algebra of F : this fact is essentially Banach’s fixed-point theorem for functions that are contractive with coefficient $1/2$.

where

$$h_n^m = \begin{cases} id_{A_m}, & \text{if } n = m, \\ g_n \circ g_{n+1} \circ \cdots \circ g_{m-1}, & \text{if } n < m, \\ f_{n-1} \circ f_{n-2} \circ \cdots \circ f_m, & \text{if } n > m. \end{cases}$$

It is easy to see that these morphisms indeed constitute a cone: in the case $n > m$ one uses that $g_{n-1} \circ f_{n-1} = id$. Hence there exists a unique morphism $i_m : A_m \rightarrow A$ such that $j_n \circ i_m = h_n^m$ for all n . In particular, $j_m \circ i_m = h_m^m = id_{A_m}$, as required in the statement of the lemma.

We must also show that $i_{m+1} \circ f_m = i_m$. By the defining property of i_m , it suffices to show that $i_{m+1} \circ f_m$ is also a cone morphism in the sense that $j_n \circ (i_{m+1} \circ f_m) = h_n^m$ for all n . And indeed, $j_n \circ i_{m+1} \circ f_m = h_n^{m+1} \circ f_m = h_n^m$ by the defining property of i_{m+1} and the definition of h_n^{m+1} .

It remains to show that $\lim_{n \rightarrow \infty} d(i_n \circ j_n, id_A) = 0$, or equivalently, that $\lim_{n \rightarrow \infty} i_n \circ j_n = id_A$ in the metric space $\mathcal{C}(A, A)$. To do so, we first show that $(i_n \circ j_n)_{n \in \omega}$ is a Cauchy sequence. Given $\epsilon > 0$, choose N large enough that $d(f_n \circ g_n, id_{A_{n+1}}) \leq \epsilon$ for all $n \geq N$. Then for all $n \geq N$,

$$\begin{aligned} d(i_n \circ j_n, i_{n+1} \circ j_{n+1}) &= d((i_{n+1} \circ f_n) \circ (g_n \circ j_{n+1}), i_{n+1} \circ j_{n+1}) \\ &= d(i_{n+1} \circ (f_n \circ g_n) \circ j_{n+1}, i_{n+1} \circ id_{A_{n+1}} \circ j_{n+1}) \\ &\leq \max(d(i_{n+1}, i_{n+1}), d(f_n \circ g_n, id_{A_{n+1}}), d(j_{n+1}, j_{n+1})) \\ &= d(f_n \circ g_n, id_{A_{n+1}}) \\ &\leq \epsilon, \end{aligned}$$

where we have used that the composition functions of an M -category are required to be non-expansive. From the ultrametric inequality one now easily obtains that $d(i_n \circ j_n, i_m \circ j_m) \leq \epsilon$ for all $n, m \geq N$. Hence $(i_n \circ j_n)_{n \in \omega}$ is a Cauchy sequence.

Since $\mathcal{C}(A, A)$ is a complete metric space, the Cauchy sequence $(i_n \circ j_n)_{n \in \omega}$ has a limit $\lim_{n \rightarrow \infty} i_n \circ j_n$. It remains to show that this limit is in fact the identity morphism on A . To do so, we show that $\lim_{n \rightarrow \infty} i_n \circ j_n$ is a cone morphism from the limiting cone $(A, j_m)_{m \in \omega}$ to itself: for all m ,

$$\begin{aligned} j_m \circ \left(\lim_{n \rightarrow \infty} i_n \circ j_n \right) &= j_m \circ \left(\lim_{n \geq m} i_n \circ j_n \right) \\ &= \lim_{n \geq m} (j_m \circ i_n \circ j_n) && \text{('o' non-expansive)} \\ &= \lim_{n \geq m} (h_m^n \circ j_n) && \text{(defining property of } i_n) \\ &= \lim_{n \geq m} j_m && \text{(by definition of } h_m^n) \\ &= j_m. \end{aligned}$$

In the second line we have again used that the composition functions of an M -category are non-expansive, hence continuous, and the fact that continuous functions preserve (metric-space) limits. We conclude that $\lim_{n \rightarrow \infty} i_n \circ j_n$ is a cone morphism from a limiting cone to itself, and therefore that $\lim_{n \rightarrow \infty} i_n \circ j_n = id_A$.

(2) implies (1): Now assume that we have a commuting diagram

$$\begin{array}{ccccccc} & & & A & & & \\ & & & \uparrow & & \swarrow & \\ & & & i_1 & & \dots & \\ & & & \vdots & & \dots & \\ & & & \vdots & & \dots & \\ & & & \vdots & & \dots & \\ A_0 & \xrightarrow{f_0} & A_1 & \xrightarrow{f_1} & \dots & \xrightarrow{f_{n-1}} & A_n \xrightarrow{f_n} \dots, \\ & \nearrow i_0 & & & & & \nwarrow i_n \end{array}$$

such that $j_n \circ i_n = id_{A_n}$ for all n and $\lim_{n \rightarrow \infty} d(i_n \circ j_n, id_A) = 0$. We must show that $(A, j_n)_{n \in \omega}$ is a limiting cone for the given ω^{op} -chain. So let $(B, b_n)_{n \in \omega}$

another cone:

$$\begin{array}{ccccccc}
 & & B & & & & \\
 & \swarrow & \downarrow & \searrow & \dots & \searrow & \\
 & b_0 & b_1 & & \dots & b_n & \dots \\
 A_0 & \xleftarrow{g_0} & A_1 & \xleftarrow{g_1} & \dots & \xleftarrow{g_{n-1}} & A_n \xleftarrow{g_n} \dots
 \end{array}$$

We aim to define a mediating morphism $q : B \rightarrow A$ as the limit of the sequence $(i_n \circ b_n)_{n \in \omega}$. We first show that this is a Cauchy sequence. The argument is completely similar to the one for the sequence $(i_n \circ j_n)_{n \in \omega}$ above: given $\epsilon > 0$, choose N large enough that $d(f_n \circ g_n, id_{A_{n+1}}) \leq \epsilon$ for all $n \geq N$. Then for all $n \geq N$,

$$\begin{aligned}
 d(i_n \circ b_n, i_{n+1} \circ b_{n+1}) &= d((i_{n+1} \circ f_n) \circ (g_n \circ b_{n+1}), i_{n+1} \circ b_{n+1}) \\
 &\leq d(f_n \circ g_n, id_{A_{n+1}}) \\
 &\leq \epsilon,
 \end{aligned}$$

and it follows that $(i_n \circ j_n)_{n \in \omega}$ is a Cauchy sequence.

Since the metric space $\mathcal{C}(B, A)$ is complete, the Cauchy sequence above has a limit: define $q = \lim_{n \rightarrow \infty} i_n \circ b_n$. We must show that q is the unique mediating morphism from the cone $(B, b_m)_{m \in \omega}$ to the cone $(A, j_m)_{m \in \omega}$. Again, the argument is as above: first,

$$\begin{aligned}
 j_m \circ q &= j_m \circ \left(\lim_{n \rightarrow \infty} i_n \circ b_n \right) = j_m \circ \left(\lim_{n \geq m} i_n \circ b_n \right) \\
 &= \lim_{n \geq m} (j_m \circ i_n \circ b_n) \\
 &= \lim_{n \geq m} (h_m^n \circ b_n) \\
 &= \lim_{n \geq m} b_m \\
 &= b_m,
 \end{aligned}$$

so q is indeed a cone morphism. Second, given another such cone morphism $r : B \rightarrow A$,

$$r = id_A \circ r = \left(\lim_{n \rightarrow \infty} i_n \circ j_n \right) \circ r = \lim_{n \rightarrow \infty} (i_n \circ j_n \circ r) = \lim_{n \rightarrow \infty} (i_n \circ b_n) = q,$$

so q is unique. We conclude that $(A, j_n)_{n \in \omega}$ is a limiting cone for the ω^{op} -chain $(A_n, g_n)_{n \in \omega}$. \square

Although not strictly necessary for our purposes, it is natural to ask whether the cocone described in Condition 2 of the lemma must be colimiting. We now show that this is the case by exploiting the generality of M -categories: the fact that \mathcal{C}^{op} is also an M -category allows for a simple proof of a limit-colimit coincidence (cf. Smyth and Plotkin [23]).

Proposition 3.3. Let \mathcal{C} be an M -category, let $(A_n, f_n, g_n)_{n \in \omega}$ be an increasing Cauchy tower in \mathcal{C} (as above), and let A be an object of \mathcal{C} . The following three conditions are equivalent:

1. A is a limit of the ω^{op} -chain $(A_n, g_n)_{n \in \omega}$.
2. A is a colimit of the ω -chain $(A_n, f_n)_{n \in \omega}$.
3. There exist a cone $(j_n)_{n \in \omega}$ from A to $(A_n, g_n)_{n \in \omega}$ and a cocone $(i_n)_{n \in \omega}$ from $(A_n, f_n)_{n \in \omega}$ to A satisfying that $j_n \circ i_n = id_{A_n}$ for all n and in addition that $\lim_{n \rightarrow \infty} d(i_n \circ j_n, id_A) = 0$.

Furthermore, in any pair consisting of a cone and a cocone that together satisfy the requirements in the third condition, the cone is limiting and the cocone is colimiting.

Proof. Lemma 3.2 shows that (1) and (3) are equivalent and that any cone $(j_n)_{n \in \omega}$ as in the third condition is limiting. The lemma also shows that these facts hold for the increasing Cauchy tower $(A_n, g_n, f_n)_{n \in \omega}$ in the M -category \mathcal{C}^{op} . But by duality, this means exactly that (2) and (3) are equivalent and that any cocone $(i_n)_{n \in \omega}$ as in the third condition is colimiting. \square

We now turn to the main result.

Theorem 3.4. Assume that the M -category \mathcal{C} satisfies any of the following (successively weaker) conditions:

1. \mathcal{C} is complete.
2. \mathcal{C} has a terminal object and limits of ω^{op} -chains.
3. \mathcal{C} has a terminal object and limits of ω^{op} -chains of split epis.
4. \mathcal{C} has a terminal object and inverse limits of increasing Cauchy towers.

Then every locally contractive functor $F : \mathcal{C}^{\text{op}} \times \mathcal{C} \rightarrow \mathcal{C}$ on \mathcal{C} has a unique fixed point up to isomorphism.

Proof. Uniqueness follows from the previous theorem, so it is enough to show that there exists some A such that $F(A, A) \cong A$. Assume that \mathcal{C} satisfies Condition 4 above and let 1 be a terminal object of \mathcal{C} . By induction on n we construct a diagram

$$A_0 \begin{array}{c} \xrightarrow{f_0} \\ \xleftarrow{g_0} \end{array} A_1 \begin{array}{c} \xrightarrow{f_1} \\ \xleftarrow{g_1} \end{array} \cdots \begin{array}{c} \xrightarrow{f_{n-1}} \\ \xleftarrow{g_{n-1}} \end{array} A_n \begin{array}{c} \xrightarrow{f_n} \\ \xleftarrow{g_n} \end{array} \cdots$$

as follows: $A_0 = 1$ and $A_{n+1} = F(A_n, A_n)$ for $n > 0$. We take g_0 to be the unique morphism from A_1 to 1 and f_0 to be an arbitrary morphism in the other direction; recall that all hom-sets in an M -category are non-empty. Finally, $f_{n+1} = F(g_n, f_n)$ and $g_{n+1} = F(f_n, g_n)$ for $n > 0$.

We now show by induction on n that this diagram is an increasing Cauchy tower. More specifically, let $c < 1$ be a contractiveness factor of F . Then, for all n :

1. $g_n \circ f_n = id_{A_n}$
2. $d(f_n \circ g_n, id_{A_{n+1}}) \leq c^n$.

For $n = 0$, Part 1 follows from the fact that $g_0 \circ f_0$ must be the identity morphism on the terminal object A_0 . Also, all distances in the spaces we consider are at most $c^0 = 1$, so Part 2 holds trivially.

As for the inductive case,

$$\begin{aligned} g_{n+1} \circ f_{n+1} &= F(f_n, g_n) \circ F(g_n, f_n) \\ &= F(g_n \circ f_n, g_n \circ f_n) \\ &= F(id_{A_n}, id_{A_n}) \quad (\text{ind. hyp.}) \\ &= id_{A_{n+1}}, \end{aligned}$$

and furthermore,

$$\begin{aligned} d(f_{n+1} \circ g_{n+1}, id_{A_{n+2}}) &= d(F(g_n, f_n) \circ F(f_n, g_n), id_{A_{n+2}}) \\ &= d(F(f_n \circ g_n, f_n \circ g_n), id_{A_{n+2}}) \\ &= d(F(f_n \circ g_n, f_n \circ g_n), F(id_{A_{n+1}}, id_{A_{n+1}})) \\ &\leq c \cdot \max(d(f_n \circ g_n, id_{A_{n+1}}), d(f_n \circ g_n, id_{A_{n+1}})) \\ &\leq c \cdot c^n \quad (\text{ind. hyp.}) \\ &= c^{n+1}, \end{aligned}$$

so both parts hold. We conclude that the diagram above is indeed an increasing Cauchy tower.

By assumption on \mathcal{C} there exists an inverse limit of this Cauchy tower, i.e., a limiting cone

$$\begin{array}{ccccccc}
 & & A & & & & \\
 & \swarrow j_0 & \downarrow j_1 & \cdots & \searrow j_n & \cdots & \\
 A_0 & \xleftarrow{g_0} & A_1 & \xleftarrow{g_1} & \cdots & \xleftarrow{g_{n-1}} & A_n & \xleftarrow{g_n} & \cdots
 \end{array}$$

By Lemma 3.2 there exist morphisms $i_n : A_n \rightarrow A$ such that $(A, i_n)_{n \in \omega}$ is a cocone from the ω -chain $(A_n, f_n)_{n \in \omega}$ to A , and such that $j_n \circ i_n = id_{A_n}$ for all n and $\lim_{n \rightarrow \infty} d(i_n \circ j_n, id_A) = 0$. In particular we have a diagram

$$\begin{array}{ccccccc}
 & & A & & & & \\
 & \swarrow i_0 & \downarrow j_1 & \cdots & \searrow j_n & \cdots & \\
 A_0 & \xleftarrow{f_0} & A_1 & \xleftarrow{f_1} & \cdots & \xleftarrow{f_{n-1}} & A_n & \xleftarrow{f_n} & \cdots \\
 & \swarrow j_0 & \downarrow i_1 & \cdots & \searrow i_n & \cdots & \\
 & & A & & & &
 \end{array}$$

which commutes in the sense that $g_n \circ j_{n+1} = j_n$ and $i_{n+1} \circ f_n = i_n$ for all n . Removing the first object of the Cauchy tower $(A_n, f_n, g_n)_{n \in \omega}$ clearly gives a new Cauchy tower, and it is easy to see that the collection of arrows i_n and j_n with $n > 0$ satisfies Condition 2 of Lemma 3.2 with respect to that Cauchy tower. Hence by that lemma, A is also a limit of the ω^{op} -chain $(A_n, g_n)_{n > 0}$ that starts from A_1 .

We now show that $F(A, A)$ is also a limit of the ω^{op} -chain $(A_n, g_n)_{n > 0}$. From that it follows that $F(A, A) \cong A$ and we are done. First we apply F to the diagram above, obtaining a diagram

$$\begin{array}{ccccccc}
 & & F(A, A) & & & & \\
 & \swarrow i'_0 & \downarrow j'_1 & \cdots & \searrow j'_n & \cdots & \\
 A_1 & \xleftarrow{f'_0} & A_2 & \xleftarrow{f'_1} & \cdots & \xleftarrow{f'_{n-1}} & A_{n+1} & \xleftarrow{f'_n} & \cdots \\
 & \swarrow j'_0 & \downarrow i'_1 & \cdots & \searrow i'_n & \cdots & \\
 & & F(A, A) & & & &
 \end{array}$$

that commutes in the same sense. Here $i'_n = F(j_n, i_n)$ and $j'_n = F(i_n, j_n)$, and similarly for the f'_n and g'_n . But by definition of the original Cauchy tower, the bottom line of the above diagram is exactly that Cauchy tower starting from A_1 . Now, by functoriality we have $j'_n \circ i'_n = F(i_n, j_n) \circ F(j_n, i_n) = F(j_n \circ i_n, j_n \circ i_n) = F(id, id) = id$ for each n , and furthermore,

$$\begin{aligned}
 \lim_{n \rightarrow \infty} d(i'_n \circ j'_n, id_{F(A, A)}) &= \lim_{n \rightarrow \infty} d(F(j_n, i_n) \circ F(i_n, j_n), id_{F(A, A)}) \\
 &= \lim_{n \rightarrow \infty} d(F(i_n \circ j_n, i_n \circ j_n), F(id_A, id_A)) \\
 &\leq \lim_{n \rightarrow \infty} c \cdot d(i_n \circ j_n, id_A) \\
 &= c \cdot \lim_{n \rightarrow \infty} d(i_n \circ j_n, id_A) \\
 &= 0,
 \end{aligned}$$

since F is contractive with factor c . Hence the morphisms in the diagram above satisfy Condition 2 of Lemma 3.2 with respect to the increasing Cauchy tower starting from A_1 . By that lemma, $F(A, A)$ is therefore a limit of the ω^{op} -chain $(A_n, g_n)_{n > 0}$. Since A is also such a limit we conclude that $F(A, A) \cong A$. \square

By the fact that \mathcal{C}^{op} is also an M -category we additionally obtain a dual version of Theorem 3.4. For example, if \mathcal{C} has an initial object and colimits of ω -chains of split monos (‘embeddings’), then every locally contractive mixed-variance functor on \mathcal{C} has a unique fixed point up to isomorphism. In the applications we have considered these dual conditions seem less useful since colimits in the categories involved are harder to describe than limits.

4 Locally compact subcategories of M -categories

The condition in Theorem 3.4 that involves Cauchy towers is included in order to accommodate categories where the hom-sets are *compact* ultrametric spaces. The simplest example is the full subcategory KBUlt_{ne} of CBUlt_{ne} consisting of compact, non-empty metric spaces. This category does not have all limits of ω^{op} -chains, not even of those chains where the morphisms are split epi. One can construct a counterexample as follows: for each $n \in \omega$, let A_n be the set $\{0, 1, \dots, n-1\}$ equipped with the discrete metric. Let $f_n : A_n \rightarrow A_{n+1}$ be the inclusion and let $g_n : A_{n+1} \rightarrow A_n$ be the function that maps n to $n-1$ and every other number to itself. We claim that the ω^{op} -chain $(A_n, g_n)_{n \in \omega}$ in KBUlt_{ne} does not have a limit. To see this, assume that $(A, j_n)_{n \in \omega}$ is a limiting cone with $j_n : A \rightarrow A_n$ for all n . By the argument in the beginning of the proof of Lemma 3.2 there exist morphisms $i_n : A_n \rightarrow A$ such that $j_n \circ i_n = \text{id}_{A_n}$ for all n . Since every A_n can in this way be embedded in A , we conclude that A contains arbitrarily large discrete subspaces. But then A cannot be totally bounded: for $\epsilon = 1/2$ there is no finite set of points such that the open balls with centers in those points and radius ϵ cover A . Hence A is not compact, a contradiction. This argument also works for KBUlt instead of KBUlt_{ne} .

The subcategory KBUlt_{ne} is merely the simplest example of a full, ‘locally compact’ subcategory of an M -category. The setting of M -categories allows for a proof that such a subcategory always inherits fixed points of functors from the full category:

Theorem 4.1. Assume that \mathcal{C} is an M -category with a terminal object and limits of ω^{op} -chains of split epis. Let I be an arbitrary object of \mathcal{C} , and let \mathcal{D} be the full subcategory of \mathcal{C} consisting of the objects A such that $\mathcal{C}(I, A)$ is a compact metric space. \mathcal{D} is an M -category with limits of increasing Cauchy towers, and hence every locally contractive functor $F : \mathcal{D}^{\text{op}} \times \mathcal{D} \rightarrow \mathcal{D}$ has a unique fixed point up to isomorphism.

Notice that the theorem refers to functors on \mathcal{D} , not on \mathcal{C} . There is in general no guarantee that a functor on \mathcal{C} restricts to one on the subcategory \mathcal{D} , and hence formulating a recursive equation by means of a functor on \mathcal{D} can require additional work [3]. In that sense, one might say that it is not exactly the problem of solving recursive equations which has been reduced to the case for the full category \mathcal{C} , but rather the problem of finding fixed-points of functors.

Proof. First, \mathcal{D} is an M -category, being a subcategory of an M -category. Second, \mathcal{D} contains each terminal object 1 of \mathcal{C} since $\mathcal{C}(I, 1)$ is the one-point metric space which is clearly compact.

We next show that \mathcal{D} has limits of increasing Cauchy towers; it then follows from Theorem 3.4 that \mathcal{D} has fixed points of locally contractive functors. To that end, let $(A_n, f_n, g_n)_{n \in \omega}$ be an increasing Cauchy-tower in \mathcal{D} (and hence also in \mathcal{C}). Each g_n is split epi, so by assumption the ω^{op} -chain $(A_n, g_n)_{n \in \omega}$ has a limiting cone $(A, (j_n)_{n \in \omega})$ in \mathcal{C} . Since \mathcal{D} is a full subcategory, it now suffices to show that the limit object A belongs to \mathcal{D} , i.e., that $\mathcal{C}(I, A)$ is compact.

Here we use the characterization of compactness from Section 2: we already know that $\mathcal{C}(I, A)$ is complete, so it remains to show that $\mathcal{C}(I, A)$ is totally bounded. First, by Lemma 3.2 applied to \mathcal{C} and the limiting cone $(A, (j_n)_{n \in \omega})$, there exists a family of morphisms $(i_n : A_n \rightarrow A)_{n \in \omega}$ satisfying certain conditions: in particular, $j_n \circ i_n = id_{A_n}$ for each n and $\lim_{n \rightarrow \infty} d(i_n \circ j_n, id_A) = 0$. Now we show that $\mathcal{C}(I, A)$ is totally bounded. Given $\epsilon > 0$, choose n large enough that $d(i_n \circ j_n, id_A) < \epsilon$. Since $\mathcal{C}(I, A_n)$ is compact, it is totally bounded. Hence there exists a finite set S of elements of $\mathcal{C}(I, A_n)$ such that for every $f \in \mathcal{C}(I, A_n)$ there is an $s \in S$ with $d(f, s) < \epsilon$. Let T be the finite subset $\{i_n \circ s \mid s \in S\}$ of $\mathcal{C}(I, A)$. Now let a be an arbitrary element of $\mathcal{C}(I, A)$. We show that a has distance less than ϵ to some element of T ; hence $\mathcal{C}(I, A)$ is totally bounded. Indeed, choose $s \in S$ such that $d(j_n \circ a, s) < \epsilon$. Then by the ultrametric inequality and the assumption that composition is non-expansive,

$$\begin{aligned} d(a, i_n \circ s) &\leq \max(d(a, i_n \circ j_n \circ a), d(i_n \circ j_n \circ a, i_n \circ s)) \\ &\leq \max(d(id_A, i_n \circ j_n), d(j_n \circ a, s)) \\ &< \epsilon. \end{aligned} \quad \square$$

Here one obtains KBUlt_{ne} by taking $\mathcal{C} = \text{CBUlt}_{\text{ne}}$ and $I = 1$. In general, for a monoidal closed \mathcal{C} , the tensor unit is an appropriate choice of I . Since we show in the next section that CBUlt_{ne} has limits of ω^{op} -chains of split epis, the theorem in particular gives:

Corollary 4.2 ([10]). Every locally contractive functor from $\text{KBUlt}_{\text{ne}}^{\text{op}} \times \text{KBUlt}_{\text{ne}}$ to KBUlt_{ne} has a unique fixed point up to isomorphism.

Moreover, the proof of the theorem above essentially works by using the hom-functor $\mathcal{C}(I, -) : \mathcal{D} \rightarrow \text{KBUlt}_{\text{ne}}$ to reduce the general case to the special case considered in the corollary.

5 Examples of categories admitting solutions

We now turn to some examples of categories that satisfy the different completeness requirements in Theorem 3.4. This section thereby illustrates which of the requirements in that theorem one might attempt to show given a particular M -category.

5.1 CBUlt_*

Consider first the category CBUlt_* of *pointed*, complete, 1-bounded ultrametric spaces. Objects are pairs (A, x) where A is a complete, 1-bounded ultrametric space and x is an element of A (a distinguished ‘point’). Morphisms from (A_1, x_1) to (A_2, x_2) are non-expansive maps f from A_1 to A_2 which ‘preserve the point’, i.e., satisfy that $f(x_1) = x_2$. We equip the hom-sets of CBUlt_* with the ‘sup’-metric, as given by the exponential in CBUlt :

$$d(f, g) = \sup\{d_{A_2}(f(x), g(x)) \mid x \in A_1\}.$$

Proposition 5.1. CBUlt_* is a complete M -category.

Proof. First, it is easy to see that CBUlt_* is an M -category. Each hom-set is non-empty since it contains the constant function whose value is the distinguished point of the codomain. The distance functions above clearly turn each hom-set into a 1-bounded ultrametric space; indeed, a sub-space of an exponential in CBUlt . Each such space is complete since the limit of a sequence of point-preserving functions is

also point-preserving. The composition functions are non-expansive since they are restrictions of composition functions from \mathbf{CBUlt} .

To see that \mathbf{CBUlt}_* is complete, it is easy to construct products and equalizers directly (as in \mathbf{CBUlt}). More abstractly, \mathbf{CBUlt}_* is the comma category $(1 \downarrow \mathbf{CBUlt})$, and the forgetful functor from this category to \mathbf{CBUlt} creates limits [17, Exercise V.1.1]. \square

5.2 $\mathbf{CBUlt}_{\text{ne}}$

We have already observed that the category $\mathbf{CBUlt}_{\text{ne}}$ of *non-empty*, complete, 1-bounded ultrametric spaces is an M -category with distance functions on hom-sets given as for exponentials. However, unlike \mathbf{CBUlt}_* , it is not complete and does not even have all limits of ω^{op} -chains. To see this, let T be a rooted tree that contains nodes of arbitrarily large depth but contains no infinite path (by König's Lemma such a tree must be infinitely branching.) For each n , let A_n be the set of nodes of T of depth n , equipped with the discrete metric. Let $g_n : A_{n+1} \rightarrow A_n$ map each node to its parent. Then the ω^{op} -chain $(A_n, g_n)_{n \in \omega}$ in $\mathbf{CBUlt}_{\text{ne}}$ does not have a limit. Indeed, a limit in $\mathbf{CBUlt}_{\text{ne}}$ would also be a limit in \mathbf{CBUlt} . But the limit of $(A_n, g_n)_{n \in \omega}$ in the complete category \mathbf{CBUlt} is the set of tuples $\{(a_n)_{n \in \omega} \mid \forall n. g_n(a_{n+1}) = a_n\}$ with the product metric; this set is empty since T does not contain any infinite path, and hence the limit does not belong to $\mathbf{CBUlt}_{\text{ne}}$.

Proposition 5.2. $\mathbf{CBUlt}_{\text{ne}}$ is an M -category with limits of ω^{op} -chains of split epis.

Proof. Since $\mathbf{CBUlt}_{\text{ne}}$ is a full subcategory of the complete category \mathbf{CBUlt} , it suffices to show that \mathbf{CBUlt} -limits of ω^{op} -chains of split epis in $\mathbf{CBUlt}_{\text{ne}}$ are non-empty. Let $(A_n, g_n)_{n \in \omega}$ be such an ω^{op} -chain, and let for each n the function f_n be a right inverse of g_n . A concrete limit in \mathbf{CBUlt} of such a chain is, as mentioned above, the set of tuples $\{(a_n)_{n \in \omega} \mid \forall n. g_n(a_{n+1}) = a_n\}$ with the product metric. Now let a_0 be an arbitrary element of A_0 (which is non-empty by assumption). It is easy to see that the limit above contains the tuple $((f_{n-1} \circ \dots \circ f_0)(a_0))_{n \in \omega}$ and is therefore also non-empty. \square

5.3 \mathbf{CBUlt}

The category \mathbf{CBUlt} is not an M -category since the set of morphisms from any non-empty metric space to the empty metric space is empty. Nevertheless, there is an obvious definition of 'locally contractive' for functors on this category, and given a locally contractive functor $F : \mathbf{CBUlt}^{\text{op}} \times \mathbf{CBUlt} \rightarrow \mathbf{CBUlt}$ that restricts to $\mathbf{CBUlt}_{\text{ne}}$, one can use the main theorem with the category $\mathbf{CBUlt}_{\text{ne}}$ to find a fixed point of F . It is not hard to see that F restricts to $\mathbf{CBUlt}_{\text{ne}}$ if and only if $F(1, 1)$ is non-empty (where 1 is the one-point metric space):

Theorem 5.3. Let $F : \mathbf{CBUlt}^{\text{op}} \times \mathbf{CBUlt} \rightarrow \mathbf{CBUlt}$ be a locally contractive functor satisfying that $F(1, 1) \neq \emptyset$. There exists a unique (up to isomorphism) non-empty $A \in \mathbf{CBUlt}$ such that $F(A, A) \cong A$.

Proof. We show that F restricts to the full subcategory $\mathbf{CBUlt}_{\text{ne}}$: given non-empty A and B , we must show that $F(A, B)$ is non-empty. Since B is non-empty there exist morphisms $f : A \rightarrow 1$ and $g : 1 \rightarrow B$ in \mathbf{CBUlt} . Then $F(f, g)$ is a function from $F(1, 1)$ to $F(A, B)$. Since $F(1, 1)$ is non-empty by assumption, the existence of such a function implies that $F(A, B)$ is non-empty too. The theorem now immediately follows from Theorem 3.4 applied to the M -category $\mathbf{CBUlt}_{\text{ne}}$. \square

Note that uniqueness is only among non-empty metric spaces: a functor F as in the theorem might furthermore satisfy that $F(\emptyset, \emptyset) = \emptyset$.

5.4 PreCBUlt_{ne}

The examples in Section 7 use the category PreCBUlt_{ne} of *pre-ordered*, non-empty, complete, 1-bounded ultrametric spaces. Objects of this category are pairs (A, \leq) consisting of an object A of CBUlt_{ne} and a preorder \leq on the underlying set of A such that the following condition holds: if $(a_n)_{n \in \omega}$ and $(b_n)_{n \in \omega}$ are converging sequences in A with $a_n \leq b_n$ for all n , then also $\lim_{n \rightarrow \infty} a_n \leq \lim_{n \rightarrow \infty} b_n$. The morphisms of the category are the non-expansive and monotone functions between such objects. We equip the hom-sets of PreCBUlt_{ne} with the usual ‘sup’-metric.

Proposition 5.4. PreCBUlt_{ne} is an M -category with limits of ω^{op} -chains of split epis.

Proof. To see that PreCBUlt_{ne} is an M -category we proceed as in the proof of Proposition 5.2. The only new thing to show is that if $(f_n)_{n \in \omega}$ is a converging sequence of monotone and non-expansive functions between objects (A, \leq_A) and (B, \leq_B) of PreCBUlt_{ne}, then the limit f is a monotone (as well as non-expansive) function. But this follows immediately from the requirement above: if $a \leq_A a'$, then

$$f(a) = \lim_{n \rightarrow \infty} f_n(a) \leq_B \lim_{n \rightarrow \infty} f_n(a') = f(a').$$

It remains to show that PreCBUlt_{ne} has limits of ω^{op} -chains of split epis. Let $(A_n, g_n)_{n \in \omega}$ be such a chain. It is easy to verify that the limit is the set of tuples $\{(a_n)_{n \in \omega} \mid \forall n. g_n(a_{n+1}) = a_n\}$ with the product metric and the product preorder, and that this set is non-empty as in the proof of Proposition 5.2. \square

5.5 Locally compact subcategories

We have already seen, using Theorem 4.1, that the full subcategory KBUlt_{ne} of CBUlt_{ne} has unique fixed points of locally contractive functors. Similarly, that theorem applied to the M -categories CBUlt_{*} and PreCBUlt_{ne} of the previous examples gives unique fixed points of locally contractive functors on the ‘compact’ variants of these two categories. Notice that for CBUlt_{*}, the choice $I = 1$ in Theorem 4.1 does not work: one must instead choose I to be the metric space consisting of two points with distance 1. (CBUlt_{*} is not cartesian closed, but it is symmetric monoidal closed with this I as tensor unit.)

6 An alternative existence theorem

We next consider an alternative existence theorem for solutions of recursive equations in M -categories. Roughly put, the overall picture is as follows. In Section 3 above we generalized the results of America and Rutten [6] to M -categories; this was done in the style of Smyth and Plotkin [23]. In this section, we outline a similar generalization of the results of Alessi et al. [4]. The resulting existence theorem can, at least informally, be viewed as a closer categorical analogy to Banach’s fixed-point theorem than the existence theorem in Section 3. In particular it will not be required that the M -category has a terminal object: any object will suffice to start the inductive construction of the solution. On the other hand, the M -category must satisfy a stronger completeness property. We do not know any applications that depend on these slightly different conditions on the category.

Let \mathcal{C} be an arbitrary M -category. In the existence proof in Section 3 we worked extensively with pairs of morphisms (f, g) such that $f : A \rightarrow B$ and $g : B \rightarrow A$ for some objects A and B of \mathcal{C} and such that $g \circ f = id_A$. Following Alessi et al. [3]

we call such pairs *embedding-projection* pairs.¹ The proof essentially takes place in a category that has such pairs as morphisms; this is made precise in, e.g., America and Rutten [6]. The alternative approach explored in this section does not depend on the embedding condition $g \circ f = id_A$ and so works with all pairs of morphisms with opposite domain and codomain. These pairs were introduced independently as ϵ -adjoint pairs in Rutten [19] and as ϵ -isometries in Alessi et al. [3]. In Alessi et al. [4] it was shown that the standard existence theorem on non-empty, complete, 1-bounded metric spaces from America and Rutten [6] could be obtained using ϵ -adjoint pairs instead of embedding-projection pairs. Here we outline a generalization of that result to M -categories.

Definition 6.1. The category \mathcal{C}^\approx has the objects of \mathcal{C} and morphisms $\iota : A \rightarrow B$ that are pairs of morphisms $\iota = \langle i, j \rangle$ of \mathcal{C} such that $i : A \rightarrow B$ and $j : B \rightarrow A$. Composition of $\iota_1 = \langle i_1, j_1 \rangle : A \rightarrow B$ and $\iota_2 = \langle i_2, j_2 \rangle : B \rightarrow C$ is defined naturally by $\iota_2 \circ \iota_1 = \langle i_2 \circ i_1, j_1 \circ j_2 \rangle : A \rightarrow C$. The identity morphism on the object A is $\langle id_A, id_A \rangle$.

We immediately remark that two objects are isomorphic in \mathcal{C} if and only if they are isomorphic in \mathcal{C}^\approx and hence we shall purposely blur the distinction. We may 'flip' any morphism $\iota = \langle i, j \rangle : A \rightarrow B$ by swapping the components to obtain a morphism $\bar{\iota} = \langle j, i \rangle : B \rightarrow A$.

Definition 6.2. The *noise* of a morphism $\iota = \langle i, j \rangle : A \rightarrow B$ in \mathcal{C}^\approx is defined as

$$\delta(\iota) = \max(d_{\mathcal{C}(A,A)}(id_A, j \circ i), d_{\mathcal{C}(B,B)}(i \circ j, id_B)).$$

Note that we rely on the M -category structure on \mathcal{C} to define the noise but make no attempt to make an M -category out of \mathcal{C}^\approx .

Intuitively, the noise measures 'how far' A and B are from each other by ι . Having $\delta(\iota) = 0$ obviously implies $j \circ i = id_A$ and $i \circ j = id_B$; in particular two objects are isomorphic if and only if there is a zero-noise morphism from one to the other. Also by definition $\delta(\iota) = \delta(\bar{\iota})$ for any morphism ι of \mathcal{C}^\approx . These two observations are somewhat analogous to the first and second of the defining axioms of an (ultra)metric space; the following lemma provides a cousin to the ultrametric inequality:

Lemma 6.3 (Noise Lemma). For $\iota_1 : A \rightarrow B$ and $\iota_2 : B \rightarrow C$ we have $\delta(\iota_2 \circ \iota_1) \leq \max(\delta(\iota_2), \delta(\iota_1))$.

Proof. Write $\iota_1 = \langle i_1, j_1 \rangle$ and $\iota_2 = \langle i_2, j_2 \rangle$. Then:

$$\begin{aligned} \delta(\iota_2 \circ \iota_1) &= \delta(\langle i_2 \circ i_1, j_1 \circ j_2 \rangle) \\ &= \max(d(id_A, j_1 \circ j_2 \circ i_2 \circ i_1), d(i_2 \circ i_1 \circ j_1 \circ j_2, id_C)) \\ &\leq \max(\max(d(id_A, j_1 \circ i_1), d(j_1 \circ id_B \circ i_1, j_1 \circ j_2 \circ i_2 \circ i_1)), \\ &\quad \max(d(i_2 \circ i_1 \circ j_1 \circ j_2, i_2 \circ id_B \circ j_2), d(i_2 \circ j_2, id_C))) \\ &\leq \max(\max(\delta(\iota_1), \delta(\iota_2)), \max(\delta(\iota_1), \delta(\iota_2))) \\ &= \max(\delta(\iota_1), \delta(\iota_2)). \end{aligned}$$

Here we have used the ultrametric inequality as well as the ubiquitous fact that the composition functions of an M -category are non-expansive. \square

In a metric space, to prove two elements equal it suffices to show that their distance is smaller than every $\epsilon > 0$. The corresponding technique in our metric-inspired setting is the following:

¹We do not, however, use any analogue of the 'projection' condition $f \circ g \sqsubseteq id_B$ from the domain-theoretic case.

Lemma 6.4 (Proximity Lemma). Two objects A and B of \mathcal{C} are isomorphic if there is a sequence of \mathcal{C}^\approx -morphisms $(\langle i_n, j_n \rangle)_{n \in \omega}$ with $\langle i_n, j_n \rangle : A \rightarrow B$ such that $\lim_{n \rightarrow \infty} \delta(\langle i_n, j_n \rangle) = 0$ and such that $(i_n)_{n \in \omega}$ and $(j_n)_{n \in \omega}$ are Cauchy sequences in $\mathcal{C}(A, B)$ and $\mathcal{C}(B, A)$, respectively.

Proof. By completeness of $\mathcal{C}(A, B)$ and $\mathcal{C}(B, A)$ we know that $\lim_{n \rightarrow \infty} i_n : A \rightarrow B$ and $\lim_{n \rightarrow \infty} j_n : B \rightarrow A$ exist. We now have:

$$\begin{aligned} d_{\mathcal{C}(A,A)} \left(id_A, \lim_{n \rightarrow \infty} j_n \circ \lim_{n \rightarrow \infty} i_n \right) &= d_{\mathcal{C}(A,A)} \left(id_A, \lim_{n \rightarrow \infty} j_n \circ i_n \right) \\ &= \lim_{n \rightarrow \infty} d_{\mathcal{C}(A,A)}(id_A, j_n \circ i_n) \\ &\leq \lim_{n \rightarrow \infty} \delta(\iota_n) \\ &= 0 \end{aligned}$$

Here we have used non-expansiveness of composition and the fact that, for any ultrametric space (X, d) , the distance function $d : X \times X \rightarrow \mathbb{R}$ is itself non-expansive and hence preserves limits. We conclude that $id_A = \lim_{n \rightarrow \infty} j_n \circ \lim_{n \rightarrow \infty} i_n$ and by symmetry we get the other way round. \square

By analogy with the standard metric argument one might try to do away with the second demand that the component sequences be Cauchy. However, as observed in Remark 4.4 of Alessi et al. [4], this is not possible. A consequence is that a ‘proper’ distance between two objects defined as the infimum of the noises of morphisms from one to the other gives only a pseudo-metric; that is, the distance between two distinct objects can be zero. This problem is explored in Section 4 of Alessi et al. [4] and solved by restricting to compact metric spaces.

Definition 6.5. A *tower* in \mathcal{C}^\approx is a sequence of pairs of objects and morphisms (A_n, ι_n) such that $\iota_n : A_n \rightarrow A_{n+1}$ for all $n \in \omega$. It is *Cauchy* if $\lim_{n \rightarrow \infty} \delta(\iota_n) = 0$, i.e., if

$$\forall \epsilon > 0. \exists N \in \mathbb{N}. \forall n \geq N. \delta(\iota_n) < \epsilon.$$

Notice that a Cauchy tower $(A_n, \iota_n)_{n \in \omega}$ where all the ι_n are embedding-projection pairs is exactly an ‘increasing Cauchy tower’ as defined in Section 3.

As in the case of standard Cauchy sequences, the objects of a Cauchy tower intuitively get arbitrarily close, measured here by the noises of the morphisms. By the Noise Lemma, i.e., due to our ultrametric setup, we immediately have that the above criterion is equivalent to one that may look more familiar:

$$\forall \epsilon > 0. \exists N \in \mathbb{N}. \forall m > n \geq N. \delta(\iota_{m-1} \circ \cdots \circ \iota_n) < \epsilon.$$

Definition 6.6. A *limit* of a Cauchy tower $(A_n, \iota_n)_{n \in \omega}$ is a pair $(A, (\gamma_n)_{n \in \omega})$ of an object and a sequence of morphisms $\gamma_n : A_n \rightarrow A$ in \mathcal{C}^\approx such that


$$\begin{array}{ccc} & A & \\ \gamma_n \nearrow & & \nwarrow \gamma_{n+1} \\ A_n & \xrightarrow{\iota_n} & A_{n+1} \end{array}$$

commutes for all $n \in \omega$ and such that $\lim_{n \rightarrow \infty} \delta(\gamma_n) = 0$.

(Proposition 6.8 below relates limits of Cauchy towers in the sense above to inverse limits of the kind considered in Section 3.)

Proposition 6.7. For any two limits $(A, (\gamma_n)_{n \in \omega})$ and $(A', (\gamma'_n)_{n \in \omega})$ of the same Cauchy tower $(A_n, \iota_n)_{n \in \omega}$ the objects A and A' are isomorphic.

Proof. This comes down to applying the Proximity Lemma. The setup is this:



We have $\gamma'_n \circ \overline{\gamma_n} : A \rightarrow A'$ for all $n \in \omega$, and since $\lim_n \delta(\gamma_n) = \lim_n \delta(\overline{\gamma_n}) = 0 = \lim_n \delta(\gamma'_n)$ we get $\lim_n \delta(\gamma'_n \circ \overline{\gamma_n}) = 0$ by the Noise Lemma. Now write $\gamma_n = \langle g_n, h_n \rangle$ and $\gamma'_n = \langle g'_n, h'_n \rangle$ for all $n \in \omega$. It remains to show that $(g'_n \circ h_n)_{n \in \omega}$ and $(g_n \circ h'_n)_{n \in \omega}$ are Cauchy sequences in the metric spaces $\mathcal{C}(A, A')$ and $\mathcal{C}(A', A)$, respectively. For any $n \in \omega$

$$\begin{aligned}
d_{\mathcal{C}(A, A')}(g'_n \circ h_n, g'_{n+1} \circ h_{n+1}) &= d_{\mathcal{C}(A, A')}(g'_{n+1} \circ i_n \circ j_n \circ h_{n+1}, g'_{n+1} \circ h_{n+1}) \\
&= d_{\mathcal{C}(A_{n+1}, A_{n+1})}(i_n \circ j_n, id_{A_{n+1}}) \\
&\leq \delta(\iota_n)
\end{aligned}$$

where we write $\iota_n = \langle i_n, j_n \rangle$. But then $(g'_n \circ h_n)_{n \in \omega}$ is Cauchy because $(A_n, \iota_n)_{n \in \omega}$ is a Cauchy tower. By symmetry $(g_n \circ h'_n)_{n \in \omega}$ is also Cauchy. \square

Notice how we use our ability to flip a morphism $\iota : A \rightarrow B$ to obtain $\bar{\iota} : B \rightarrow A$; in the category of embedding-projection pairs this is not possible in general.

We say that \mathcal{C}^\approx is *tower-complete* if all Cauchy towers have limits. Verifying this condition directly may be an arduous task. The following criterion is sufficient:

Proposition 6.8. \mathcal{C}^\approx is tower-complete if \mathcal{C} has inverse limits of Cauchy towers.

We omit the proof. The arguments follow those in the first part of the proof of Lemma 3.2, but are more involved since we no longer restrict to embedding-projection pairs. More specifically, the cone $(h_n^m)_{n \in \omega}$ from A_m to $(A_n, g_n)_{n \in \omega}$ in that proof must now be defined as follows:

$$h_n^m = \begin{cases} k_n, & \text{if } n = m, \\ g_n \circ g_{n+1} \circ \cdots \circ g_{m-1} \circ k_m, & \text{if } n < m, \\ k_n \circ f_{n-1} \circ f_{n-2} \circ \cdots \circ f_m, & \text{if } n > m. \end{cases}$$

where each $k_n : A_n \rightarrow A_n$ is obtained as a limit of a Cauchy sequence:

$$k_n = \lim_{p \geq n} (g_n \circ g_{n+1} \circ \cdots \circ g_{p-1} \circ f_p \circ f_{p-1} \circ \cdots \circ f_n).$$

For a domain-theoretic analogue of dropping the restriction to embedding-projection pairs, see Taylor [24].

6.1 Fixed points of Functors

We now move on to apply the theory to build fixed points of functors. We say that a functor $\Phi : \mathcal{C}^\approx \rightarrow \mathcal{C}^\approx$ is *contractive* if there is a $c < 1$ such that $\delta(\Phi(\iota)) \leq c \cdot \delta(\iota)$ holds for all morphisms ι of \mathcal{C}^\approx . Similarly it is called *non-expansive* if the noises do not increase, i.e., if $\delta(\Phi(\iota)) \leq \delta(\iota)$ holds for all ι . We may build functors on \mathcal{C}^\approx from functors on \mathcal{C} :

Proposition 6.9. Let $F : \mathcal{C}^{\text{op}} \times \mathcal{C} \rightarrow \mathcal{C}$. We define $F^\approx : \mathcal{C}^\approx \rightarrow \mathcal{C}^\approx$ by

$$F^\approx(A) = F(A, A), \quad F^\approx(\langle i, j \rangle) = \langle F(j, i), F(i, j) \rangle$$

for any object A and any morphism $\iota = \langle i, j \rangle : A \rightarrow B$ of \mathcal{C}^\approx . This constitutes a well defined functor. Moreover, if F is locally contractive then F^\approx is contractive and if F is locally non-expansive then F^\approx is non-expansive.

We saw in Theorem 3.1 that a locally contractive functor $F : \mathcal{C}^{\text{op}} \times \mathcal{C} \rightarrow \mathcal{C}$ has at most one fixed point up to isomorphism (for arbitrary \mathcal{C} .) One can give an alternative proof of that fact using the Proximity Lemma and the contractive functor F^\approx derived from F . But even in the concrete case $\mathcal{C} = \text{CBUlt}_{\text{ne}}$ it is an open question whether every contractive endofunctor on \mathcal{C}^\approx has at most one fixed point [3, p. 7].

Just as non-expansive maps between metric spaces are continuous and thus preserve limits of sequences, we have the following proposition as an immediate consequence of the above definitions:

Proposition 6.10. Non-expansive functors preserve limits of Cauchy towers. That is, for any non-expansive functor $\Phi : \mathcal{C}^\approx \rightarrow \mathcal{C}^\approx$ and any Cauchy tower $(A_n, \iota_n)_{n \in \omega}$ with limit $(A, (\gamma_n)_{n \in \omega})$ we have that $(\Phi(A_n), \Phi(\iota_n))_{n \in \omega}$ is a Cauchy tower with limit $(\Phi(A), (\Phi(\gamma_n))_{n \in \omega})$.

Theorem 6.11. If \mathcal{C}^\approx is nonempty and tower-complete then any contractive functor $\Phi : \mathcal{C}^\approx \rightarrow \mathcal{C}^\approx$ has a fixed point, i.e., an object A of \mathcal{C}^\approx with $A \cong \Phi(A)$.

Proof. Much of the theory above targets this proof; it is quite short and analogous to the proof of Banach's fixed-point theorem.

Let A_0 be any object of \mathcal{C}^\approx and define $A_{n+1} = \Phi(A_n)$ for every $n \in \omega$. Let $\iota_0 : A_0 \rightarrow A_1$ be any morphism of \mathcal{C}^\approx and define $\iota_{n+1} = \Phi(\iota_n) : A_{n+1} \rightarrow A_{n+2}$ for every $n \in \omega$. We can always initiate this process: \mathcal{C}^\approx was assumed to have an object, and $\iota_0 : A_0 \rightarrow A_1$ always exists as the hom-sets of an M -category are non-empty.

It is immediate by the contractiveness of Φ that $(A_n, \iota_n)_{n \in \omega}$ is a Cauchy tower and hence has a limit $(A, (\gamma_n)_{n \in \omega})$ as \mathcal{C}^\approx was assumed tower-complete. A contractive functor is in particular non-expansive and non-expansive functors preserves limits of Cauchy towers, so $(\Phi(A_n), \Phi(\iota_n))_{n \in \omega} = (A_{n+1}, \iota_{n+1})_{n \in \omega}$ is a Cauchy tower too with limit $(\Phi(A), (\Phi(\gamma_n))_{n \in \omega})$. But $(A, (\gamma_{n+1})_{n \in \omega})$ is a limit of $(A_{n+1}, \iota_{n+1})_{n \in \omega}$ too and uniqueness of limits (Proposition 6.7) gives $A \cong \Phi(A)$. \square

Combining Proposition 6.8, Proposition 6.9, and Theorem 6.11 we have:

Theorem 6.12. If \mathcal{C} has an object and has inverse limits of Cauchy towers then every locally contractive functor $F : \mathcal{C}^{\text{op}} \times \mathcal{C} \rightarrow \mathcal{C}$ has a unique fixed point up to isomorphism.

Notice that here we require *all* Cauchy towers to have inverse limits, not just the increasing ones. Therefore Theorem 6.12 does not immediately imply Theorem 3.4.

7 Applications

This section contains a series of examples of recursive equations motivated by recent and ongoing work in semantics. In all but the first of the examples we do not consider exactly those equations that arise from applications; for clarity we consider simplified variants that capture the essence of the circularity issues. We conclude the section by discussing in what sense the generality of M -categories is needed in applications.

7.1 Realizability semantics of dynamically allocated store

The first two examples of recursive equations come from realizability semantics of dynamically allocated store. In recent work [9] the authors presented a model that allows for simple parametricity-style reasoning about imperative abstract data types in an ML-like language with universal types, recursive types, and reference types. As in Standard ML, references are dynamically allocated during program execution.

Here is a brief outline of the model. First, the model is based on a realizability interpretation [5] over a certain recursively defined predomain V . In addition, we follow earlier work on modeling simple integer references [8] and use a Kripke-style possible worlds model. Here, however, the set of worlds needs to be recursively defined since we treat general references. Semantically, a world maps locations to semantic types, which, following the general realizability idea, are certain world-indexed families of relations on V : this introduces a circularity between semantic types and worlds that precludes a direct definition of either. Thus we are led to solving recursive (metric-space) equations of approximately the following form

$$\begin{aligned}\mathcal{W} &\cong \mathbb{N} \rightarrow_{fin} \mathcal{T} \\ \mathcal{T} &\cong \mathcal{W} \rightarrow_{mon} CUREl(V)\end{aligned}$$

(see below) even in order to define the space in which types will be modeled.

We now describe these equations in more detail. $CUREl(V)$ is the set of binary relations on V that satisfy certain technical requirements. The metric on $CUREl(V)$ is defined essentially as in earlier work on realizability semantics [5], using the fact that V is a canonical solution to a predomain equation. The space $\mathbb{N} \rightarrow_{fin} \mathcal{T}$ consists of partial functions from \mathbb{N} to \mathcal{T} with finite domain: the distance between two functions with different domains is 1, while the distance between two functions with the same domain is given as a maximum of pointwise distances. The space $\mathbb{N} \rightarrow_{fin} \mathcal{T}$ (and hence also \mathcal{W}) is equipped with an extension order: for $\Delta, \Delta' \in \mathbb{N} \rightarrow_{fin} \mathcal{T}$ we take $\Delta \leq \Delta'$ to mean that $\text{dom}(\Delta) \subseteq \text{dom}(\Delta')$ and that $\Delta(n) = \Delta'(n)$ for all n in $\text{dom}(\Delta)$. Finally, in order to ensure soundness of the interpretation, we require the usual ‘Kripke monotonicity’: the space $\mathcal{W} \rightarrow_{mon} CUREl(V)$ should consist of functions that are both non-expansive and monotone with respect to the extension order on \mathcal{W} and the inclusion order on $CUREl(V)$.

In order to apply the main theorem to solve these equations, we have to express them in terms of a mixed-variance functor on an M -category. There are two approaches. First, one can ‘solve for worlds’ by defining a contravariant functor F on PreCBUlt_{ne} such that

$$F(X, \leq) = (\mathbb{N} \rightarrow_{fin} \frac{1}{2} ((X, \leq) \rightarrow_{mon} CUREl(V))), \leq'$$

where \leq' is the extension order on partial functions, as defined above. (Here the $\frac{1}{2}$ is needed in order to ensure that F is locally contractive.) Then (\mathcal{W}, \leq) can be defined as the unique fixed point of F .

Alternatively, one can ‘solve for types’ [9] by defining a contravariant functor G on CBUlt_{ne} (or on CBUlt as in Section 5.3) such that

$$G(X) = \frac{1}{2} ((\mathbb{N} \rightarrow_{fin} X) \rightarrow_{mon} CUREl(V)).$$

Then \mathcal{T} can be defined as the unique fixed point of G . In this case we do not use the generality of M -categories: instead we exploit that the two mutually recursive equations above have a form that allows one to solve them in CBUlt_{ne} by combining them into a single recursive equation in the right way. In the next example such an approach will not be possible; there, M -categories seem to be needed.

Remark. Even though metric spaces appear naturally in this example through earlier work on realizability, one might ask whether the equations above cannot instead be solved in a category of domains. Indeed they can, but the solution does not appear to be useful for our purpose. The main reason is that the quantitative information given by the metric-space approach appears to be needed in order to model reference types [9].

7.2 A more advanced model of store

In the previous example, semantic types were modeled as world-indexed families of binary relations on the predomain V of ‘untyped values’. The intuitive idea is that worlds provide information about the currently allocated references, and that the interpretation of a type grows as more references are allocated.

The fact that relations are binary allows one to use the model to prove equivalences between programs that allocate references dynamically. However, the set of worlds $\mathcal{W} \cong \mathbb{N} \rightarrow_{fin} \mathcal{T}$ of the previous example only allows for fairly limited equivalence proofs. There, a world is no more than a single ‘semantic store typing’ that only allows one to describe situations where the two programs under consideration allocate references in lockstep.

Ongoing work suggests that the metric-space approach allows one to solve an equation involving more advanced Kripke worlds in the style of Ahmed et al. [2], and thereby allows for more advanced reasoning about local state.² Here we present a simplified equation that illustrates the main circularity issue. Let $S = \mathbb{N} \rightarrow_{fin} V$ be the set of *stores*, i.e., partial maps with finite domain from \mathbb{N} to V . Whereas the simple worlds of the previous example induce a binary relation on stores that require two related stores to have the same domain, we now seek an alternative definition of worlds that induce a more liberal relation on stores.

The intuitive idea is that a world \mathcal{W}' consists of a finite sequence of ‘islands’ I [2], each of which induces a local requirement on stores by describing how two specific parts of two given stores are required to be related. Consider the metric-space equations

$$\begin{aligned} \mathcal{W}' &\cong I^* \\ I &\cong \sum_{N_1, N_2 \in P_{fin}(\mathbb{N})} \frac{1}{2} (\mathcal{W}' \rightarrow_{mon} CURel(S)_{N_1, N_2}) \end{aligned}$$

which are to be understood as follows. The space $CURel(S)$ [9] consists of binary relations on stores satisfying certain technical conditions; it is equipped with a metric in the same way as $CURel(V)$ above. Given finite subsets N_1 and N_2 of \mathbb{N} , the sub-space $CURel(S)_{N_1, N_2}$ of $CURel(S)$ only contains relations with *support* (N_1, N_2) , i.e., $R \in CURel(S)_{N_1, N_2}$ and $(s_1, s_2) \in R$ implies $(s'_1, s'_2) \in R$ if $s_1(n) = s'_1(n)$ for all n in N_1 and $s_2(n) = s'_2(n)$ for all n in N_2 . Intuitively, such relations are local in the sense that they only depend on locations from N_1 and N_2 , respectively. The sum on the right-hand side of the second equation consists of triples (N_1, N_2, f) ; the distance between two such triples is 1 if either of the first two components differ, and the distance between the third components otherwise. Finally, assuming that the second equation holds, the space I^* consists of finite sequences of triples (N_1, N_2, f) such that the first components are pairwise disjoint, and similarly for the second components: the ‘islands’ must not overlap. The extension order on I^* , and hence on \mathcal{W}' , is sequence containment; the maps of the second equation are monotone with respect to this order and the inclusion order on $CURel(S)_{N_1, N_2}$.

²Ahmed et al. do not solve the recursive equation they consider, but instead work with a family of sets that are, intuitively, approximations to a solution.

Because of the different dependencies on finite subsets of integers, it does not seem possible to combine the two equations above into one equation in CBUlt_{ne} : some extra structure on metric spaces is needed, no matter what one tries to ‘solve for’. Indeed, the equations can be ‘solved for worlds’ by defining a contravariant functor H on $\text{PreCBUlt}_{\text{ne}}$ directly from the equations,

$$H(X, \leq) = \left(\left(\sum_{N_1, N_2 \in P_{\text{fin}}(\mathbb{N})} \frac{1}{2} ((X, \leq) \rightarrow_{\text{mon}} \text{CUREl}_{N_1, N_2}(S)) \right)^*, \leq' \right)$$

where \leq' is the extension order on sequences, and then letting \mathcal{W}' be the unique fixed point of H .

7.3 Storable locks

In recent work by Gotsman et al. separation logic has been extended to reason about storable locks and threads [15]. As observed in *loc. cit.* the natural model of predicates involves a circular definition because locks protecting invariants (predicates) can be stored in the heap. However, Gotsman et al. side-step this issue by restricting the storable locks to protect only a statically determined finite set of kinds of invariants. In ongoing work, Birkedal and Buisse are generalizing the work by Gotsman et al. by solving a suitable recursive equation. The equation is

$$\text{UPred} \cong \frac{1}{2} \left((\mathbb{N} \rightarrow_{\text{fin}} (\mathbb{N} + (\mathbb{N} \times \text{UPred}))) \rightarrow_{\text{mon}} P\downarrow(\mathbb{N}) \right).$$

Here $P\downarrow(\mathbb{N})$ is the complete, bounded ultrametric space consisting of downwards-closed subsets of \mathbb{N} ; this set forms a complete Heyting algebra. The idea is that a semantic predicate is a $P\downarrow(\mathbb{N})$ -valued predicate on heaps, which are maps from locations (numbers) to either numbers or pairs (k, I) consisting of a thread id k and a semantic predicate I . The latter is used if the location is a lock, held by thread k and protecting the invariant I .

This equation can be solved in CBUlt_{ne} by solving for UPred (much as in the example in Section 7.1), or by solving for heaps by defining a contravariant functor on $\text{PreCBUlt}_{\text{ne}}$.

7.4 Semantics of nested Hoare triples

In recent work, Schwinghammer et al. [21] investigate the semantics of separation logic for higher-order store. Their uniform admissible subsets of heaps form the basic building block when interpreting the assertions of the logic. Since assertions in general depend on invariants for stored code (because of higher-order store), the space of semantic predicates consists of functions $\mathcal{W} \rightarrow \text{UAdm}$ from a set of ‘worlds,’ describing the invariants, to the collection of uniform admissible subsets of heaps. The set UAdm is an ultrametric space with metric given as for $\text{CUREl}(V)$ above. But, the invariants for stored code are themselves semantic predicates, and hence the space of worlds \mathcal{W} should be ‘the same’ as $\mathcal{W} \rightarrow \text{UAdm}$. Thus the following equation is solved in CBUlt_{ne} :

$$\mathcal{W} \cong \frac{1}{2} (\mathcal{W} \rightarrow \text{UAdm}).$$

7.5 Discussion

As we have seen, three of the four examples above could be treated by solving recursive equations in CBUlt_{ne} , i.e., without using the generality of M -categories. The fourth example, the advanced model of store in Section 7.2, does seem to require

M -categories; however, a recent simplification [14] leads to an equation that can be solved in $\mathbf{CBUlt}_{\text{ne}}$.

In a little more detail, the general situation is the following. Instead of directly finding a fixed point of an endofunctor $F : \mathcal{C} \rightarrow \mathcal{C}$ on an M -category, one can attempt to express the functor as $F = G \circ H$ for functors $H : \mathcal{C} \rightarrow \mathbf{CBUlt}_{\text{ne}}$ and $G : \mathbf{CBUlt}_{\text{ne}} \rightarrow \mathcal{C}$, and then find a fixed point of $H \circ G : \mathbf{CBUlt}_{\text{ne}} \rightarrow \mathbf{CBUlt}_{\text{ne}}$ instead of F . In this way, only the classical existence theorem is needed. (It could of course also be that one came up with $H \circ G$ instead of F in the first place.) In the first example above, this is the approach that could be used to ‘solve for types’ instead of ‘solving for worlds.’

Despite of this possibility of rewriting a recursive equation to an equation in $\mathbf{CBUlt}_{\text{ne}}$, we find M -categories genuinely useful in applications for the following reasons. First, there is no guarantee that one can rewrite a given recursive equation; as already noted, we cannot see how to treat the example in Section 7.2 by solving an equation in $\mathbf{CBUlt}_{\text{ne}}$. Second, and more subtly: in the first and third examples above, the ‘rewriting’ approach still depends on the fact that $\mathbf{PreCBUlt}_{\text{ne}}$ is an M -category. More specifically, to obtain a well-defined endofunctor on $\mathbf{CBUlt}_{\text{ne}}$ in these examples, one implicitly uses the defining property of objects of $\mathbf{PreCBUlt}_{\text{ne}}$ that guarantees that hom-sets of $\mathbf{PreCBUlt}_{\text{ne}}$ are *complete* metric spaces. Without having identified $\mathbf{PreCBUlt}_{\text{ne}}$ as an M -category, this leads to some ad-hoc calculations and results about monotone function spaces, as can be seen in Birkedal et al. [9].

8 Domain equations: from O -categories to M -categories

As another illustration of M -categories, we present a general construction that gives for every O -category \mathcal{C} (see below) a derived M -category \mathcal{D} . In addition, the construction gives for every locally continuous mixed-variance functor F on \mathcal{C} a locally contractive mixed-variance functor G on \mathcal{D} such that a fixed point of G , necessarily unique by Theorem 3.1, is the same as a fixed point of F that furthermore satisfies the ‘minimal invariance’ condition of Pitts [18].³ Thus, generalized domain equations can be solved in M -categories.

The construction generalizes and improves an earlier one due to Baier and Majster-Cederbaum (BM) [7] which is for the particular category \mathbf{Cppo}_{\perp} of pointed cpos and strict, continuous functions (or full subcategories thereof.) More precisely, taking \mathcal{C} to be a full subcategory of \mathbf{Cppo}_{\perp} in our Proposition 8.2 below gives a result that strengthens Lemma 4.18 of BM. In general, the goal of that earlier work is to relate recursive domain equations over full subcategories of \mathbf{Cppo}_{\perp} to recursive equations over full subcategories of a particular category \mathbf{CMS} of complete metric spaces. Working with those particular categories instead of arbitrary O -categories and M -categories complicates the relations one can obtain: for example, Theorem 3 of BM only applies to a restricted class of domain equations that does not include general function spaces. The reason is that the construction of BM, which must be applied to a full subcategory of \mathbf{Cppo}_{\perp} , does not yield a category that is (in any obvious way) a full subcategory of \mathbf{CMS} . It is, however, an M -category in which every locally contractive functor has a unique fixed point. We hence believe to have at least partially answered the question left open in the conclusion of BM whether a suitable notion of correspondence exists for general domain equations.

Rank-ordered cpos [7], recently re-discovered under the name ‘uniform cpos’ [9], arise from a particular instance of an M -category obtained from the construction

³The latter is in turn the same as a bifree algebra for F in the same sense as in Theorem 3.1. See the argument in Pitts [18].

here, namely by taking $\mathcal{C} = \mathbf{Cppo}_\perp$. The extra metric information in that category, as compared with the underlying O -category, is useful in realizability models [1, 5]; in particular it is used to define the metric on the set of binary relations $CURel(V)$ in the work described in Section 7.1 [9]. In earlier work, Abadi and Plotkin [1, Section 8] give a metric-space formulation of a realizability model of polymorphism and recursive types. They note that the extra metric information can be used to model subtypes and bounded quantification.

We now turn to the results. An O -category [23] is a category \mathcal{C} where each hom-set $\mathcal{C}(A, B)$ is equipped with an ω -complete partial order, usually written \sqsubseteq , and where each composition function is continuous with respect to these orders. A functor $F : \mathcal{C}^{\text{op}} \times \mathcal{C} \rightarrow \mathcal{C}$ is *locally continuous* if each function on hom-sets that it induces is continuous.

Assume now that \mathcal{C} is an O -category such that each hom-set $\mathcal{C}(A, B)$ contains a least element $\perp_{A,B}$ and such that the composition functions of \mathcal{C} are strict: $f \circ \perp_{A,B} = \perp_{A,C} = \perp_{B,C} \circ g$ for all f and g . We construct an M -category \mathcal{D} of ‘rank-ordered \mathcal{C} -objects’ as follows. An object $(A, (\pi_n)_{n \in \omega})$ of \mathcal{D} is a pair consisting of an object A of \mathcal{C} and a family of endomorphisms $\pi_n : A \rightarrow A$ in \mathcal{C} that satisfies the following four requirements:

- (1) $\pi_0 = \perp_{A,A}$.
- (2) $\pi_m \sqsubseteq \pi_n$ for all $m \leq n$.
- (3) $\pi_m \circ \pi_n = \pi_n \circ \pi_m = \pi_{\min(m,n)}$ for all m and n .
- (4) $\bigsqcup_{n \in \omega} \pi_n = id_A$.

(See also the rank-ordered sets of Baier and Majster-Cederbaum [7] and the projection spaces of de Vries [12].) Then, a morphism from $(A, (\pi_n)_{n \in \omega})$ to $(A', (\pi'_n)_{n \in \omega})$ in \mathcal{D} is a morphism f from A to A' in \mathcal{C} that is *uniform* [1] in the sense that $\pi'_n \circ f = f \circ \pi_n$ for all n . Composition and identities in \mathcal{D} are the same as in \mathcal{C} . Finally, the distance function on a hom-set $\mathcal{D}((A, (\pi_n)_{n \in \omega}), (A', (\pi'_n)_{n \in \omega}))$ is defined as follows:

$$d(f, g) = \begin{cases} 2^{-\max\{n \in \omega \mid \pi'_n \circ f = \pi'_n \circ g\}} & \text{if } f \neq g \\ 0 & \text{if } f = g. \end{cases}$$

To see that d is well-defined, suppose that $f \neq g$. Then there must exist a greatest number n such that $\pi'_n \circ f = \pi'_n \circ g$. Indeed, $n = 0$ is such a number by (1) above and strictness of the composition functions of \mathcal{C} . If the equation holds for arbitrarily large n , then by (3) above it holds for all n . But then by (4) above and the fact that the composition functions of \mathcal{C} are continuous,

$$f = id_{A'} \circ f = \left(\bigsqcup_{n \in \omega} \pi'_n \right) \circ f = \bigsqcup_{n \in \omega} (\pi'_n \circ f) = \bigsqcup_{n \in \omega} (\pi'_n \circ g) = \dots = g,$$

a contradiction. Hence d is well-defined.

Proposition 8.1. \mathcal{D} is an M -category.

Proof. First, each hom-set $\mathcal{D}((A, (\pi_n)_{n \in \omega}), (A', (\pi'_n)_{n \in \omega}))$ is non-empty: it contains the element $\perp_{A,A'}$ since $\pi'_n \circ \perp_{A,A'} = \perp_{A,A'} \circ \pi_n = \perp_{A,A'}$ by strictness. Second, it is easy to see that the distance function on such a hom-set gives rise to a 1-bounded ultrametric space. Third, the composition functions of \mathcal{D} are non-expansive: it suffices to see that if $\pi''_n \circ f_1 = \pi''_n \circ f_2$ and $\pi'_n \circ g_1 = \pi'_n \circ g_2$, then $\pi''_n \circ (f_1 \circ g_1) = \pi''_n \circ f_2 \circ g_1 = f_2 \circ \pi'_n \circ g_1 = f_2 \circ \pi'_n \circ g_2 = \pi''_n \circ (f_2 \circ g_2)$.

It remains to show that each hom-set is a complete metric space. Let $(f_m)_{m \in \omega}$ be a Cauchy sequence. It follows from the definition of d that for each $n \in \omega$ there

exists g_n such that $\pi'_n \circ f_m = g_n$ for all sufficiently large m . Then by (2) above the sequence $(g_n)_{n \in \omega}$ is increasing: given n , we have $g_n = \pi'_n \circ f_m \sqsubseteq \pi'_{n+1} \circ f_m = g_{n+1}$ for all sufficiently large m .

The supremum $g = \sqcup_{n \in \omega} g_n$ is the limit of the sequence $(f_m)_{m \in \omega}$. Indeed, given an arbitrary number k , by (3) we have $\pi'_k \circ g_n = \pi'_k \circ \pi'_n \circ f_m = \pi'_k \circ f_m = g_k$ for all $n \geq k$ (and sufficiently large m), and therefore, by continuity of composition, $\pi'_k \circ (\sqcup_{n \in \omega} g_n) = \sqcup_{n \geq k} (\pi'_k \circ g_n) = g_k = \pi'_k \circ f_m$ for all sufficiently large m . This shows that $d(f_m, g) \leq 2^{-k}$ for all sufficiently large m . Hence g is indeed the limit of the sequence $(f_m)_{m \in \omega}$.

We must show that g is uniform. First, each g_n is uniform since for all k and all sufficiently large m we have $\pi'_k \circ g_n = \pi'_k \circ \pi'_n \circ f_m = \pi'_k \circ \pi'_n \circ f_m = \pi'_n \circ f_m \circ \pi_k = g_n \circ \pi_k$. Second, g is uniform since each g_n is: $\pi'_k \circ (\sqcup_{n \in \omega} g_n) = \sqcup_{n \in \omega} (\pi'_k \circ g_n) = \sqcup_{n \in \omega} (g_n \circ \pi_k) = (\sqcup_{n \in \omega} g_n) \circ \pi_k$. In conclusion, each hom-set is complete, and \mathcal{D} is an M -category. \square

Now let $F : \mathcal{C}^{\text{op}} \times \mathcal{C} \rightarrow \mathcal{C}$ be a locally continuous functor. We construct a locally contractive functor $G : \mathcal{D}^{\text{op}} \times \mathcal{D} \rightarrow \mathcal{D}$ from F :

- On objects, G is given by

$$G((A, (\pi_n^A)_{n \in \omega}), (B, (\pi_n^B)_{n \in \omega})) = (F(A, B), (\pi_n^{A,B})_{n \in \omega})$$

where $\pi_0^{A,B} = \perp$ and $\pi_{n+1}^{A,B} = F(\pi_n^A, \pi_n^B)$ for all n .

- On morphisms, G is the same as F , i.e., $G(f, g) = F(f, g)$.

To see that G is well-defined on objects, we must verify conditions (1)-(4) in the definition of objects of \mathcal{D} . Here (1) is immediate, (3) follows from strictness of composition and functoriality of F , and (2) and (4) follow from local continuity of F . In addition, given morphisms $f : (A', (\pi_n^{A'})_{n \in \omega}) \rightarrow (A, (\pi_n^A)_{n \in \omega})$ and $g : (B, (\pi_n^B)_{n \in \omega}) \rightarrow (B', (\pi_n^{B'})_{n \in \omega})$, we must show that $G(f, g) = F(f, g)$ is a well-defined morphism in \mathcal{D} . Clearly, $\pi_0^{A',B'} \circ F(f, g) = F(f, g) \circ \pi_0^{A,B}$, and for all n ,

$$\begin{aligned} \pi_{n+1}^{A',B'} \circ F(f, g) &= F(\pi_n^{A'}, \pi_n^{B'}) \circ F(f, g) \\ &= F(f \circ \pi_n^{A'}, \pi_n^{B'} \circ g) \\ &= F(\pi_n^A \circ f, g \circ \pi_n^B) \\ &= F(f, g) \circ F(\pi_n^A, \pi_n^B) \\ &= F(f, g) \circ \pi_{n+1}^{A,B}. \end{aligned}$$

Finally, G is locally contractive with factor $1/2$: it suffices to see that if $\pi_n^A \circ f_1 = \pi_n^A \circ f_2$ and $\pi_n^{B'} \circ g_1 = \pi_n^{B'} \circ g_2$, then

$$\begin{aligned} \pi_{n+1}^{A',B'} \circ F(f_1, g_1) &= F(\pi_n^{A'}, \pi_n^{B'}) \circ F(f_1, g_1) \\ &= F(f_1 \circ \pi_n^{A'}, \pi_n^{B'} \circ g_1) \\ &= F(\pi_n^A \circ f_1, \pi_n^{B'} \circ g_1) \\ &= F(\pi_n^A \circ f_2, \pi_n^{B'} \circ g_2) \\ &= \pi_{n+1}^{A',B'} \circ F(f_2, g_2). \end{aligned}$$

Proposition 8.2. Let G be constructed from F as above, and let A be an object of \mathcal{C} . The following two conditions are equivalent.

- (1) There exists an isomorphism $i : F(A, A) \rightarrow A$ such that

$$id_A = fix(\lambda e^{C(A,A)}. i \circ F(e, e) \circ i^{-1}).$$

(Here fix is the least-fixed-point operator.)

- (2) There exists a family of morphisms $(\pi_n)_{n \in \omega}$ such that $\bar{A} = (A, (\pi_n)_{n \in \omega})$ is the unique fixed-point of G up to isomorphism.

Proof. (1) implies (2): Define the \mathcal{C} -morphisms $\pi_n : A \rightarrow A$ by induction on n : $\pi_0 = \perp_{A,A}$ and $\pi_{n+1} = i \circ F(\pi_n, \pi_n) \circ i^{-1}$. We must show that $(A, (\pi_n)_{n \in \omega})$ is an object of \mathcal{D} by verifying the four requirements in the definition of \mathcal{D} . The first requirement is immediate by definition, the second and third requirements are easy to show by induction, and the fourth requirement is exactly the assumption that $id_A = fix(\lambda e^{C(A,A)}. i \circ F(e, e) \circ i^{-1})$. Now let $\bar{A} = (A, (\pi_n)_{n \in \omega})$. It remains to show that $G(\bar{A}, \bar{A}) \cong \bar{A}$; uniqueness follows from Theorem 3.1. We claim that the isomorphism $i : F(A, A) \rightarrow A$ in \mathcal{C} is also an isomorphism from $G(\bar{A}, \bar{A})$ to \bar{A} in \mathcal{D} , i.e., that both i and its inverse i^{-1} in \mathcal{C} are uniform with respect to the families of morphisms $(\pi_n^{A,A})_{n \in \omega}$ and $(\pi_n)_{n \in \omega}$ on $F(A, A)$ and A , respectively. Clearly $\pi_0 \circ i = i \circ \pi_0^{A,A}$ by strictness. Also,

$$\pi_{n+1} \circ i = (i \circ F(\pi_n, \pi_n) \circ i^{-1}) \circ i = i \circ F(\pi_n, \pi_n) = i \circ \pi_{n+1}^{A,A}$$

by the definitions of π_{n+1} and $\pi_{n+1}^{A,A}$. So i is uniform. The proof that i^{-1} is uniform is completely similar. In conclusion, $i : G(\bar{A}, \bar{A}) \rightarrow \bar{A}$ is an isomorphism in \mathcal{D} .

(2) implies (1): Assume that $i : G(\bar{A}, \bar{A}) \rightarrow \bar{A}$ is an isomorphism in \mathcal{D} . Then $i : F(A, A) \rightarrow A$ is clearly also an isomorphism in \mathcal{C} ; formally one applies the forgetful functor $\mathcal{D} \rightarrow \mathcal{C}$ to i . Since i is uniform, as are all morphisms in \mathcal{D} , we have that $\pi_{n+1} \circ i = i \circ \pi_{n+1}^{A,A} = i \circ F(\pi_n, \pi_n)$, and hence that $\pi_{n+1} = i \circ F(\pi_n, \pi_n) \circ i^{-1}$. By the definition of objects of \mathcal{D} we furthermore have that $\pi_0 = \perp_{A,A}$ and that $\sqcup_{n \in \omega} \pi_n = id_A$. But then

$$fix(\lambda e^{C(A,A)}. i \circ F(e, e) \circ i^{-1}) = \bigsqcup_{n \in \omega} \pi_n = id_A. \quad \square$$

It remains to discuss how completeness properties of \mathcal{C} transfer to \mathcal{D} . One can show, using the O -category variant of Lemma 3.2 [23], that the forgetful functor from \mathcal{D} to \mathcal{C} creates terminal objects and limits of ω^{op} -chains of split epis. Alternatively, by imposing an additional requirement on \mathcal{C} one can show that the forgetful functor creates *all* limits: for a given limit in \mathcal{C} , the induced bijection between cones and mediating morphisms must be an isomorphism in the category of cpos (where cones are ordered pointwise, using the order on each hom-set). That requirement is in particular satisfied by the usual concrete categories of cpos.

9 Discussion and related work

We now discuss the level of generality of our results and clarify what the contribution is compared to other results in the literature.

Metric spaces vs. ultrametric spaces. In the definition of an M -category we require that hom-sets are *ultrametric* spaces and not merely general metric spaces. In most of the related work on metric-space equations that we cite there is no such restriction; only Rutten [19] restricts to ultrametric spaces.

The main reason that we restrict to ultrametric spaces is that doing so is sufficient for the kind of applications in semantics we consider. Intuitively, the distinguishing feature of ultrametric spaces is that one can think of the distance $d(x, x')$ between two elements as specifying some ‘level’ up to which x and x' are in perfect agreement. Such distance functions are often used: Turi and Rutten [25, p.529] observed in 1998 that “in semantics, one usually works with these more specific structures [ultrametric spaces].” On the other hand, there are by now applications

in semantics where general (pseudo-)metric spaces must be used. In particular, general metric spaces can be used to capture similarity in probabilistic transition systems in a quantitative way [13, 26].

If we had considered general metric spaces, some definitions and proofs would have become more complicated. More specifically, the definition of an increasing Cauchy tower below would need to be modified in the style of the definition of a ‘converging tower’ in America and Rutten [6]. With that change, we expect that at least the results of Section 3 and 4 would carry through also for general metric spaces. A high-level explanation for the added complications is that the category of complete, general metric spaces is not cartesian closed, only symmetric monoidal closed [16].

Actually, one could specialize further than ultrametric spaces. All the ultrametric spaces we consider in the applications in Section 7 are of a special kind: all non-zero distances are of the form 2^{-n} for some natural number $n \geq 0$. In practice, this means that when working with those applications, and in particular in calculations, one considers the metric-space structure as given by a family of equivalence relations $=_n$ where $x =_n y$ holds if and only if $d(x, y) \leq 2^{-n}$. One could perhaps specialize even further, by requiring that these relations $=_n$ are obtained from ‘projection’ functions [7, 12]. We did not restrict to any such structures, but found ultrametric spaces to give simple proofs at a reasonable level of generality.

Generalized ultrametric spaces. Apart from the distinction between general metric spaces and ultrametric spaces, there are other choices one could consider in the definition of M -categories. First, it is likely that the restriction to *non-empty* metric spaces could be removed, but that this would require unpleasant special cases in the uniqueness and existence theorems.

Generalizing in a different direction than metric spaces, one could consider categories where the hom-sets are *generalized* ultrametric spaces [19]. In these spaces, only the ultrametric inequality and the axiom $d(x, x) = 0$ are required: the distance function need not be symmetric, and distinct elements can have distance 0. These spaces are particularly nice from an (enriched) category theoretic viewpoint, since they are exactly categories enriched over the preorder category $[0, 1]^{\text{op}}$. On the other hand, limits in such spaces are somewhat more complicated than in ordinary metric spaces, and some basic intuitions need to be revisited: for example, non-expansive functions are not necessarily continuous.

A remarkable aspect of generalized ultrametric spaces is that they generalize both metric spaces and preorders. We find this particularly interesting in light of the comparison between M -categories and O -categories in Section 8: one could imagine that categories enriched over (complete) generalized ultrametric spaces would allow for a result about solving recursive equations that generalizes both our work and the classical result of Smyth and Plotkin. We do not know whether there exists such a common generalization and leave it as a direction for future work.

Finally, one could perhaps go even further and generalize from the preorder category $[0, 1]^{\text{op}}$ to arbitrary quantales. Here the goal would be to prove results that relate to the work of Wagner [27] in the same way that Smyth and Plotkin’s work on O -categories relates to Scott’s classical inverse limit construction. Wagner [27, p. 92] briefly discusses such a possibility.

Contributions compared to related work. As mentioned in the introduction, the idea of considering categories with metric spaces as hom-sets has been used in earlier work [10, 20]. We now turn to a more detailed description of what is new in our results.

First of all, while the results in Sections 4 (on locally compact subcategories)

and 8 (on relating O -categories to M -categories) are generalizations of earlier work, we find these generalizations sufficiently far-reaching and non-trivial to be considered new in their own right.

As for the main results in Section 3 about solutions to recursive equations, the contribution with respect to previous work consists in Lemma 3.2, Proposition 3.3, and Theorem 3.4. First, Lemma 3.2 (and its immediate consequence, Proposition 3.3) shows a limit-colimit coincidence result for general M -categories. Theorem 4.15 of Rutten and Turi [20] gives some, but not all of this result. That theorem only refers to cones in a category \mathcal{C}^E of embedding-projection pairs⁴, and does not relate these cones to limiting cones or colimiting cocones in \mathcal{C} . In particular, the theorem does not show that an inverse limit in \mathcal{C} can be turned into a cone in a category of embedding-projection pairs. This argument, which is an essential part of the construction of fixed points of functors, is later given in the concrete case of a category of complete metric spaces only (Theorem 4.23).

Second, Theorem 3.4 is new: we do not know of any other existence theorem about fixed points of functors for a general class of ‘metric-enriched’ categories. A small contribution here is the identification of suitable completeness conditions on M -categories, together with examples that illustrate their use. In particular, the condition of having all inverse limits of split epis is needed for the prime example of an M -category, namely CBUit_{ne} itself. This condition seems to be of general use given that the hom-sets of an M -category are required to be non-empty.

Still, it is fair to say that (as one reviewer put it) the ingredients of the existence theorem already occur in various places in the literature. Perhaps a more proper way to evaluate our main result is to consider how the applications in Section 7 would be treated using these previous results, notably Theorem 4.25 of Rutten and Turi [20]. Although the recursive equations in the applications can be solved in this way, we think that the route requires sufficient extra work that it useful to state a general existence theorem that directly applies to mixed-variance functors on arbitrary M -categories. Most importantly, in the absence of our Lemma 3.2 it is, as noted above, necessary to argue ‘by hand’ that inverse limits of ω^{op} -chains in $\text{PreCBUit}_{\text{ne}}$ can be turned into cones in a category of embedding-projection pairs.

10 Conclusion

We have generalized the standard solution of recursive equations over complete ultrametric spaces [6] to the abstract setting of M -categories where, in the style of Smyth and Plotkin [23], the focus is on the metric structure on the morphisms rather than the objects. We have furthermore outlined an alternative existence theorem which is, at least informally, a closer categorical analogy to Banach’s fixed-point theorem.

We have given a general account of ‘compact’ variants of such categories, showing that these subcategories always inherit solutions of recursive equations from the full categories. As another application we have presented a construction that provides a correspondence between solutions of generalized domain equations in O -categories with solutions of equations in M -categories.

In addition, we have sketched a number of applications from denotational semantics. In particular, the application in Section 7.2 requires a solution to a recursive equation over metric spaces with additional structure; our results provide such a solution.

⁴We assume that the second \mathcal{C} in the statement of that theorem should have been \mathcal{C}^E .

Acknowledgments We thank the anonymous referees for helpful comments and suggestions. This research was partially funded by the grant ‘Modular Reasoning about Software’ from The Danish Council for Independent Research, Natural Sciences.

References

- [1] M. Abadi and G. D. Plotkin. A per model of polymorphism and recursive types. In J. Mitchell, editor, *Proceedings of the Fifth Annual IEEE Symposium on Logic in Computer Science*, pages 355–365, Philadelphia, Pennsylvania, June 1990. IEEE Computer Society Press.
- [2] A. Ahmed, D. Dreyer, and A. Rossberg. State-dependent representation independence. In Z. Shao and B. C. Pierce, editors, *Proceedings of the Thirty-Sixth Annual ACM Symposium on Principles of Programming Languages*, pages 340–353, Savannah, Georgia, Jan. 2009. ACM Press.
- [3] F. Alessi, P. Baldan, and G. Bellè. A fixed-point theorem in a category of compact metric spaces. *Theoretical Computer Science*, 146(1&2):311–320, 1995.
- [4] F. Alessi, P. Baldan, G. Bellè, and J. J. M. M. Rutten. Solutions of functorial and non-functorial metric domain equations. In S. Brookes, M. Main, A. Melton, and M. Mislove, editors, *Proceedings of the 11th Conference on the Mathematical Foundations of Programming Semantics*, volume 1 of *Electronic Notes in Theoretical Computer Science*, pages 1–12. Elsevier, 1995.
- [5] R. M. Amadio. Recursion over realizability structures. *Information and Computation*, 91(1):55–85, 1991.
- [6] P. America and J. J. M. M. Rutten. Solving reflexive domain equations in a category of complete metric spaces. *Journal of Computer and System Sciences*, 39(3):343–375, 1989.
- [7] C. Baier and M. E. Majster-Cederbaum. The connection between initial and unique solutions of domain equations in the partial order and metric approach. *Formal Aspects of Computing*, 9(4):425–445, 1997.
- [8] N. Benton and B. Leperchey. Relational reasoning in a nominal semantics for storage. In P. Urzyczyn, editor, *Typed Lambda Calculi and Applications, 7th International Conference, TLCA 2005*, number 3461 in Lecture Notes in Computer Science, pages 86–101, Nara, Japan, Apr. 2005. Springer.
- [9] L. Birkedal, K. Støvring, and J. Thamsborg. Realizability semantics of parametric polymorphism, general references, and recursive types. In L. de Alfaro, editor, *Foundations of Software Science and Computational Structures, 12th International Conference, FOSSACS 2009*, number 5504 in Lecture Notes in Computer Science, pages 456–470, York, United Kingdom, Mar. 2009. Springer.
- [10] F. van Breugel and J. Warmerdam. Solving domain equations in a category of compact metric spaces. Report CS-R9424, CWI, Amsterdam, 1994.
- [11] J. W. de Bakker and J. Zucker. Processes and the denotational semantics of concurrency. *Information and Control*, 54:70–120, 1982.
- [12] F.-J. de Vries. Projection spaces and recursive domain equations. *Information Processing Society of Japan SIG Notes*, 95(114):37–38, 1995. Extended abstract. Retrieved from <http://ci.nii.ac.jp/Detail/detail.do?LOCALID=ART0003278737&lang=en>.

- [13] J. Desharnais, V. Gupta, R. Jagadeesan, and P. Panangaden. Metrics for labelled Markov processes. *Theoretical Computer Science*, 318(3):323–354, 2004.
- [14] D. Dreyer, G. Neis, A. Rossberg, and L. Birkedal. A relational modal logic for higher-order stateful ADTs. In M. V. Hermenegildo and J. Palsberg, editors, *Proceedings of the Thirty-Seventh Annual ACM Symposium on Principles of Programming Languages*, pages 185–198, Madrid, Spain, 2010. ACM Press.
- [15] A. Gotsman, J. Berdine, B. Cook, N. Rinetzky, and M. Sagiv. Local reasoning for storable locks and threads. In Z. Shao, editor, *Programming Languages and Systems, 5th Asian Symposium, APLAS 2007*, number 4807 in Lecture Notes in Computer Science, pages 19–37, Singapore, Nov. 2007. Springer.
- [16] F. W. Lawvere. Metric spaces, generalized logic, and closed categories. In *Rendiconti del Seminario Matematico e Fisico di Milano, XLIII*. Tipografia Fusi, Pavia, 1973.
- [17] S. Mac Lane. *Categories for the Working Mathematician, Second Edition*, volume 5 of *Graduate Texts in Mathematics*. Springer, 1998.
- [18] A. M. Pitts. Relational properties of domains. *Information and Computation*, 127:66–90, 1996.
- [19] J. J. M. M. Rutten. Elements of generalized ultrametric domain theory. *Theoretical Computer Science*, 170(1-2):349–381, 1996.
- [20] J. J. M. M. Rutten and D. Turi. On the foundation of final semantics: Non-standard sets, metric spaces, partial orders. In J. W. de Bakker, W.-P. de Roever, and G. Rozenberg, editors, *Proceedings of the REX Workshop on Semantics: Foundations and Applications*, number 666 in Lecture Notes in Computer Science, pages 477–530, Beekbergen, The Netherlands, June 1992. Springer.
- [21] J. Schwinghammer, L. Birkedal, B. Reus, and H. Yang. Nested Hoare triples and frame rules for higher-order store. In E. Grädel and R. Kahle, editors, *Proceedings of the 18th EACSL Annual Conference on Computer Science Logic*, volume 5771 of *Lecture Notes in Computer Science*, pages 440–454, Coimbra, Portugal, Sept. 2009. Springer.
- [22] M. B. Smyth. Topology. In S. Abramsky, D. Gabbay, and T. S. E. Maibaum, editors, *Handbook of Logic in Computer Science*. Oxford University Press, 1992.
- [23] M. B. Smyth and G. D. Plotkin. The category-theoretic solution of recursive domain equations. *SIAM Journal on Computing*, 11(4):761–783, 1982.
- [24] P. Taylor. Homomorphisms, bilimits and saturated domains, 1987. Manuscript. Available from <http://www.paultaylor.eu/domains/>.
- [25] D. Turi and J. J. M. M. Rutten. On the foundations of final coalgebra semantics. *Mathematical Structures in Computer Science*, 8(5):481–540, 1998.
- [26] F. van Breugel and J. Worrell. A behavioural pseudometric for probabilistic transition systems. *Theoretical Computer Science*, 331(1):115–142, 2005.
- [27] K. R. Wagner. *Solving Recursive Domain Equations with Enriched Categories*. PhD thesis, Carnegie Mellon University, Pittsburgh, Pennsylvania, 1994.