Low distortion embeddings of infinite metric spaces into the real line

Stefan Geschke [1]

Department of Mathematics, Boise State University
1910 University Drive, Boise, ID 83725-1555, USA

Abstract

We present a proof of a Ramsey-type theorem for infinite metric spaces due to Matoušek. Then we show that for every $K > 1$ every uncountable Polish space has a perfect subset that K-bi-Lipschitz embeds into the real line. Finally we study decompositions of infinite separable metric spaces into subsets that, for some $K > 1$, K-bi-Lipschitz embed into the real line.

Key words: metric space, Ramsey theory, distortion

1 Introduction

For $K > 1$ we say that an embedding between metric spaces is a K-*embedding* if the embedding and its inverse are both Lipschitz of constant K. Bourgain, Figiel and Milman [2] showed that for every constant $K > 1$ every finite metric space of size n has a subspace of size $\Omega(\log n)$ that K-embeds into ℓ^2. In particular, for every m there is n such that every metric space of size n has a subspace of size m that K-embeds into ℓ^2. (See [1] for more recent results in this direction.) This theorem clearly reminds one of the finite Ramsey Theorem (see [7]) since it says that a large metric space has a large subspace on which the metric is in some sense canonical. The finite Ramsey Theorem says that for every coloring $c : [X]^2 \to \{0, 1\}$ of the two-element subsets of an n-element set X, X has a subset H of size $\Omega(\log n)$ such that H is *homogeneous*, i.e., such that c is constant on the set $[H]^2$ of two-element subsets of H.

[1] The research for this paper was supported by G.I.F. Research Grant No. I-802-195.6/2003.

The infinite Ramsey Theorem says that for every coloring $c : [X]^2 \to \{0, 1\}$ of the two-element subsets of an infinite set X there is an infinite homogeneous set $H \subseteq X$. The theorem of Bourgain, Figiel and Milman mentioned above has an infinite analog as well. An even stronger Ramsey-type theorem for infinite metric spaces has been stated, without a proof, by Matoušek [12]: For every $K > 1$, every infinite metric space has an infinite subset that K-embeds either into the real line or into an infinite uniform space, i.e., a metric space in which any two distinct points have the same distance. A finite version of this result was proved by Karloff, Rabani and Ravid [9] in the context of motion planning.

It is well known that the infinite Ramsey Theorem fails at larger cardinals. There is a coloring $c : [\mathbb{R}]^2 \to \{0, 1\}$ such that no uncountable subset of $\mathbb{R}$ is homogeneous with respect to c [7, Section 6.4, Theorem 1]. However, once some regularity condition is imposed on the coloring $c : [X]^2 \to \{0, 1\}$, there are large homogeneous sets. Galvin showed that if X is a Polish space without isolated points and $c : [X]^2 \to \{0, 1\}$ is a coloring such that the sets $\{(x, y) \in X : \{x, y\} \in c^{-1}(i)\}$, $i \in \{0, 1\}$, have the Baire property in X^2, then X has a perfect subset that is homogeneous with respect to c (see [10, Theorem 19.7]). We prove a metric analog of this Theorem: If X is an uncountable Polish space and $K > 1$, then X has a perfect subset that K-embeds into the real line. Note that no regularity assumptions are necessary in this case, simply because the metric on X is continuous by default. To simplify the notation we call a set that K-embeds into the real line *K-linear*. Having established the existence of large K-linear sets, we proceed further and show that consistently every separable metric space can be covered by a small number of K-linear sets. This development is similar to the situation in continuous Ramsey theory [5], but technically simpler. We finally come up with a strictly increasing sequence $(K_n)_{n\in\omega}$ of real numbers > 1 such that for all $n \in \omega$ it is consistent that less K_{n+1}-linear subsets of $\mathbb{R}^2$ are needed to cover the whole plane than K_n-linear subsets.

1.1 Outline of the paper

Section 2 is devoted to a detailed proof of Matoušek's theorem for infinite metric spaces, the metric analog of Galvin's theorem for uncountable Polish spaces and some compactness results that will be used in Section 3. Section 3 mainly deals with consistency results related to decompositions of Polish spaces into K-linear sets. Here some knowledge of set theory and in particular forcing is assumed. [8] and [11] are excellent sources concerning this subject.

2 Results in ZFC

2.1 Embedding sequences into $\mathbb{R}$

Definition 1 *a) Let $f : X \to Y$ be an injection between metric spaces (X, d_X) and (Y, d_Y). For a real constant $K \geq 1$, f is a* K-embedding *if for all $x, y \in X$ with $x \neq y$ we have*

$$\frac{1}{K} \leq \frac{d_Y(f(x), f(y))}{d_X(x, y)} \leq K.$$

b) A metric space (X, d) is K-linear *if it K-embeds into $\mathbb{R}$.*

Note that for K-embeddings $f : X \to Y$ and $g : Y \to Z$, $g \circ f : X \to Z$ is a K^2-embedding. Obviously, 1-embeddings are just isometric embeddings.

In this subsection we show that sequences in metric spaces are can be K-embedded into $\mathbb{R}$ if they either diverge or converge sufficiently fast. We interpolate the embeddings into $\mathbb{R}$ by embeddings into ultrametric spaces.

Definition 2 *A metric space (X, d_X) is* ultrametric *if for all $x, y, z \in X$ we have $d_X(x, z) \leq \max(d_X(x, y), d_X(y, z))$.*

Definition 3 *A sequence $(x_n)_{n\in\omega}$ is* anti-Cauchy *with respect to a metric d if for every $k \in \omega$ there is $n_0 \in \omega$ such that for all $n, m \in \omega$ with $n_0 \leq n < m$ we have $k \leq d(x_n, x_m)$.*

Lemma 4 *Let $K > 1$. Suppose $(x_n)_{n\in\omega}$ is a sequence that is anti-Cauchy with respect to some metric d. Then $X = \{x_n : n \in \omega\}$ has an infinite subset that is K-linear.*

The proof of this lemma is based on

Lemma 5 *Let $K > 1$ and $\varepsilon = 1 - \frac{1}{K}$. Suppose $(x_n)_{n\in\omega}$ is a sequence without repetitions such that, with respect to some metric d, the following holds:*

For every $n \in \omega$ and all $i, j < n$,

$$d(x_i, x_j) \leq \varepsilon \cdot d(x_i, x_n).$$

We define an ultrametric by letting

$$d_{\mathrm{u}}(x_i, x_j) = d(x_0, x_{\max(i,j)})$$

for all $i, j \in \omega$ with $i \neq j$.

Then the identity map on $\{x_n : n \in \omega\}$ is a K-embedding with respect to d and d_u. Moreover, $(\{x_n : n \in \omega\}, d_u)$ is K-linear.

PROOF. We first show that d_u is indeed an ultrametric. Observe that the sequence $(d(x_0, x_n))_{n\in\omega}$ is increasing and hence $d_u(x_i, x_k) \leq d_u(x_j, x_\ell)$ if $i < k$, $j < \ell$ and $k \leq \ell$. Now let $i, j, k \in \omega$ be pairwise distinct. If $\max(i, j, k) = j$, then $d_u(x_i, x_k) \leq d_u(x_j, x_k)$. If $\max(i, j, k) \in \{i, k\}$, then $d_u(x_i, x_k) = d_u(x_i, x_j)$ or $d_u(x_i, x_k) = d_u(x_j, x_k)$. In any case we have

$$d_u(x_i, x_k) \leq \max(d_u(x_i, x_j), d_u(x_j, x_k)).$$

In order to show that the identity map is a K-embedding with respect to d and d_u let $i, j \in \omega$ be such that $i < j$. Then

$$\frac{d_u(x_i, x_j)}{d(x_i, x_j)} = \frac{d(x_0, x_j)}{d(x_i, x_j)} \leq \frac{d(x_0, x_i) + d(x_i, x_j)}{d(x_i, x_j)} = 1 + \frac{d(x_0, x_i)}{d(x_i, x_j)} \leq 1 + \varepsilon \leq K.$$

On the other hand,

$$\frac{d_u(x_i, x_j)}{d(x_i, x_j)} = \frac{d(x_0, x_j)}{d(x_i, x_j)} \geq \frac{d(x_i, x_j) - d(x_0, x_i)}{d(x_i, x_j)} = 1 - \frac{d(x_0, x_i)}{d(x_i, x_j)} \geq 1 - \varepsilon \geq \frac{1}{K}.$$

Finally, consider the embedding

$$e : \{x_n : n \in \omega\} \to \mathbb{R}; x_n \mapsto d(x_0, x_n).$$

For $i, j \in \omega$ with $i < j$ we have

$$\frac{|e(x_i) - e(x_j)|}{d_u(x_i, x_j)} = \frac{d(x_0, x_j) - d(x_0, x_i)}{d(x_0, x_j)} \leq 1$$

and

$$\frac{|e(x_i) - e(x_j)|}{d_u(x_i, x_j)} = \frac{d(x_0, x_j) - d(x_0, x_i)}{d(x_0, x_j)} \geq 1 - \frac{d(x_0, x_i)}{d(x_0, x_j)} \geq 1 - \varepsilon \geq \frac{1}{K}.$$

This shows that e is a K-embedding with respect to d_u and the usual metric on $\mathbb{R}$. □

Proof of Lemma 4 If $(x_n)_{n\in\omega}$ is anti-Cauchy, then it can easily be thinned out to a sequence as in Lemma 5 for the constant $\sqrt{K}$. Lemma 4 now follows by the remark after Definition 1. □

Observe that a metric space X contains an anti-Cauchy sequence if and only if its set of distances is unbounded. Therefore Lemma 4 implies

Corollary 6 *Let $K > 1$. Then every metric space X with an unbounded set of distances has an infinite subset that is K-linear.*

For Cauchy sequences we have the following analog of Lemma 4:

Lemma 7 *Let $K > 1$. Suppose $(x_n)_{n\in\omega}$ is a sequence without repetitions that is Cauchy with respect to some metric d. Then $\{x_n : n \in \omega\}$ has an infinite subset that is K-linear.*

The proof of Lemma 7 uses

Lemma 8 *Let $K > 1$ and $\varepsilon = 1 - \frac{1}{K}$. Suppose $(x_n)_{n\in\omega}$ is a sequence without repetitions such that, with respect to some metric d, the following holds:*

For every $n \in \omega$ and all $i, j, k > n$,

$$d(x_i, x_j) \leq \varepsilon \cdot d(x_n, x_k).$$

We define an ultrametric by letting

$$d_{\mathrm{u}}(x_i, x_j) = \inf_{k>i} d(x_i, x_k)$$

for all $i, j \in \omega$ with $i < j$.

Then the identity map on $\{x_n : n \in \omega\}$ is a K-embedding with respect to d and d_{u}. Moreover, $(\{x_n : n \in \omega\}, d_{\mathrm{u}})$ is K-linear.

PROOF. We show that d_{u} is an ultrametric. First observe that $d_{\mathrm{u}}(x_i, x_j)$ only depends on the smaller one of the indices. Moreover, the sequence $(d_{\mathrm{u}}(x_i, x_{i+1}))_{i\in\omega}$ is decreasing since for all $j > i+1$ and all $k > i$ we have $d(x_{i+1}, x_j) \leq \varepsilon \cdot d(x_i, x_k)$ and hence

$$d_{\mathrm{u}}(x_{i+1}, x_{i+2}) = \inf_{j>i+1} d(x_{i+1}, x_j) \leq \inf_{k>i} \varepsilon \cdot d(x_i, x_k) = \varepsilon \cdot d_{\mathrm{u}}(x_i, x_{i+1}).$$

If $i, j, k \in \omega$ are pairwise distinct, then either $j = \min(i, j, k)$ or $\min(i, j, k) \in \{i, k\}$. In the first case $d_{\mathrm{u}}(x_i, x_k) \leq d_{\mathrm{u}}(x_i, x_j)$. In the second case $d_{\mathrm{u}}(x_i, x_k) = d_{\mathrm{u}}(x_i, x_j)$ or $d_{\mathrm{u}}(x_i, x_k) = d_{\mathrm{u}}(x_j, x_k)$. In any case we have

$$d_{\mathrm{u}}(x_i, x_k) \leq \max(d_{\mathrm{u}}(x_i, x_j), d_{\mathrm{u}}(x_j, x_k)).$$

Now let $i, j \in \omega$ with $i < j$. Then

$$\frac{d_{\mathrm{u}}(x_i, x_j)}{d(x_i, x_j)} = \frac{\inf_{k>i} d(x_i, x_k)}{d(x_i, x_j)} \leq 1.$$

On the other hand,

$$\frac{d_{\mathrm{u}}(x_i,x_j)}{d(x_i,x_j)} = \frac{\inf_{k>i} d(x_i,x_k)}{d(x_i,x_j)} \geq \frac{d(x_i,x_j) - \sup_{k>i} d(x_j,x_k)}{d(x_i,x_j)} \geq 1-\varepsilon = \frac{1}{K}.$$

It follows that the identity map is a K-embedding with respect to d and d_{u}.

Finally consider the embedding

$$e : \{x_n : n \in \omega\} \to \mathbb{R}; x_n \mapsto d_{\mathrm{u}}(x_n, x_{n+1}).$$

For all $i, j \in \omega$ with $i < j$ we have

$$\frac{|x_i - x_j|}{d_{\mathrm{u}}(x_i,x_j)} = \frac{d_{\mathrm{u}}(x_i,x_{i+1}) - d_{\mathrm{u}}(x_j,x_{j+1})}{d_{\mathrm{u}}(x_i,x_{i+1})} \leq 1$$

and

$$\begin{aligned}\frac{|x_i - x_j|}{d_{\mathrm{u}}(x_i,x_j)} &= \frac{d_{\mathrm{u}}(x_i,x_{i+1}) - d_{\mathrm{u}}(x_j,x_{j+1})}{d_{\mathrm{u}}(x_i,x_{i+1})} \\ &= 1 - \frac{\inf_{k>j} d(x_j,x_k)}{\inf_{\ell>i} d(x_i,x_\ell)} \geq 1 - \sup_{k>j,\ell>i} \frac{d(x_j,x_k)}{d(x_i,x_\ell)} \geq 1-\varepsilon = \frac{1}{K}.\end{aligned}$$

It follows that e is a K-embedding with respect to d_{u} and the usual metric on $\mathbb{R}$. □

Proof of Lemma 8 Since $(x_n)_{n\in\omega}$ has no repetitions, we may assume, after removing a point from the sequence, that $(x_n)_{n\in\omega}$ does not converge to any of the x_n. For each $n \in \omega$ the sequence $(d(x_n, x_i))_{i\in\omega}$ is Cauchy in $\mathbb{R}$ since $(x_n)_{n\in\omega}$ is Cauchy with respect to d. Let $d_n = \lim_{i\to\infty} d(x_n, x_i)$. Note that $d_n > 0$.

Let $\varepsilon = 1 - \frac{1}{\sqrt{K}}$. By recursion on $m \in \omega$ we choose a strictly increasing sequence $(n_m)_{m\in\omega}$ in ω such that for all $m \in \omega$ and all $i, j, k \geq n_{m+1}$ we have

$$d(x_i, x_j) \leq \frac{\varepsilon}{2} \cdot d_{n_m}$$

and

$$\frac{1}{2} \cdot d_{n_m} \leq d(x_{n_m}, x_k).$$

Now if $i, j, k, m \in \omega$ are such that $i, j, k > m$, then

$$d(x_{n_i}, x_{n_j}) \leq \frac{\varepsilon}{2} \cdot d_{n_m} \leq \varepsilon \cdot d(x_{n_m}, x_{n_k}).$$

In other words, the sequence $(x_{n_m})_{m\in\omega}$ satisfies the requirements in Lemma 8 for the constant $\sqrt{K}$. Lemma 7 now easily follows by the remark after Definition 1. □

If X is an infinite subset of $\mathbb{R}^n$, then either it is unbounded and therefore contains an anti-Cauchy sequence or its closure is compact and therefore X contains a Cauchy sequence. From Lemma 4 and Lemma 7 we now easily obtain

Corollary 9 *Let $K > 1$. Then every infinite set $X \subseteq \mathbb{R}^n$ has an infinite subset Y that is K-linear.*

2.2 Metric spaces with a set of non-zero distances that is bounded from above and from below

We use the infinite Ramsey Theorem to show that every infinite metric space that neither has distinct points of very small nor of very large distance has an infinite subsets where any two distinct points have nearly the same distance.

Theorem 10 (Ramsey, see Theorem 5 in [7, Section 1]) *Let X be an infinite set and let c be a map from the set $[X]^2$ of two-element subsets of X into a finite set of colors. Then X has an infinite subset H such that c is constant on $[H]^2$.*

Note that an easy induction suffices to get Theorem 10 from its version for two colors mentioned in the introduction.

Definition 11 *A metric space X is* uniform *if there is a constant D such that any two distinct points in X have distance D. X is K-*uniform *if it K-embeds into a uniform metric space.*

Clearly, a uniform metric space is ultrametric.

Observe that if the non-zero distances in a metric space X only vary by a factor of at most K, then X is K-uniform. Just choose any $D > 0$ that occurs as a distance in X and replace the metric on X by the uniform metric with distance D.

On the other hand, if X is K-uniform, then the non-zero distances in X only vary by a factor of at most K^2.

Lemma 12 *Let $K > 1$. Let (X, d) be an infinite metric space and assume that there are $\varepsilon > 0$ and $N \in \omega$ such that for all $x, y \in X$ with $x \neq y$ we have $\varepsilon \leq d(x, y) < N$.*

Then X has an infinite subset Y that is K-uniform.

PROOF. For every $n \in \omega$ let $c_n = \varepsilon \cdot K^n$. Let $M \in \omega$ be maximal with $c_M < N$. For all $x, y \in X$ with $x \neq y$ let $c(x, y)$ be the unique $i \in \{0, \dots, M\}$ such that $d(x, y) \in [c_i, c_{i+1})$. By the infinite Ramsey Theorem, there is an infinite set $Y \subseteq X$ such that for some $i \in \{0, \dots, M\}$ for all $x, y \in Y$ with $x \neq y$ we have $c(x, y) = i$. Now for all $a, b, x, y \in Y$ with $a \neq b$ and $x \neq y$ we have

$$c_i \leq d(a, b) < c_{i+1} = K \cdot c_i \leq K \cdot d(x, y).$$

By the remark after Definition 11, this shows that Y is K-uniform. □

It it worth pointing out that the infinite Ramsey Theorem can be easily derived from Lemma 12. If $c : [X]^2 \to \{0, 1\}$ is a coloring, we define a metric on X by letting $d(x, y) = 1$ if $x \neq y$ and $c(x, y) = 1$, and $d(x, y) = 2$ if $x \neq y$ and $c(x, y) = 0$. By Lemma 12, X has an infinite subset H on which d is uniform. If the distances are 1, all two-element subsets of H have color 1, if all distances are 2, all two-element subsets of H have color 0.

2.3 A Ramsey-type theorem for infinite metric spaces

Theorem 13 (Matoušek [12]) *Let X be an infinite metric space and $K > 1$. Then there is an infinite set $Y \subseteq X$ that is either K-linear or K-uniform.*

PROOF. Let d denote the metric on X. By Corollary 6 we may assume that the set of distances in X is bounded (from above). Fix $n > 0$. For all $x, y \in X$ with $x \neq y$ let

$$c_n(x, y) = \begin{cases} 0 & \text{if } d(x, y) < \frac{1}{n}, \\ 1 & \text{if } d(x, y) \geq \frac{1}{n} \end{cases}$$

By recursion on n we construct a decreasing sequence $(H_n)_{n \in \omega}$ of infinite subsets of X as follows:

Let $H_0 = X$. Assume we have constructed H_n. By the infinite Ramsey Theorem, H_n has an infinite subset H_{n+1} such that for some $i \in \{0, 1\}$ for all $x, y \in H_{n+1}$ with $x \neq y$ we have $c_{n+1}(x, y) = i$. We say that i is the *color* of H_{n+1}.

Observe that if for some $n > 0$ the set H_n is of color 1, then for every $m > n$ the set H_m is of color 1. If H_n is of color 0, then for every $m > n$, H_m is of color 0 or 1.

We are left with two cases.

(1) For every $n > 0$ the color of H_n is 0.

(2) There is $m > 0$ such that for all $n \geq m$ the color of H_n is 1.

In Case (1) we choose a sequence $(x_n)_{n\in\omega}$ without repetitions such that for every $n \in \omega$, $x_n \in H_n$. It is easily checked that the sequence is Cauchy. It now follows from Lemma 7 that $\{x_n : n \in \omega\}$ has an infinite subset Y that is K-linear.

In Case (2) it follows from Lemma 12 that H_m has an infinite subset Y that is K-uniform. □

Since the K-embeddings in $\mathbb{R}$ in the proof of Theorem 13 all factor through low distortion embeddings into ultrametric spaces, we actually have the following slightly more explicit theorem:

Theorem 14 *Let $K > 1$. Then every infinite metric space X has an infinite subset Y that K-embeds into an ultrametric space that is either K-linear or uniform.*

2.4 A Ramsey-type theorem for complete metric spaces.

We now prove a metric analog of Galvin's theorem about Borel colorings of the two-element subsets of a complete metric space (see [10, Section 19.B]). Recall that a set in a metric space is *perfect* if it is closed and has no isolated points. We tacitly assume that perfect sets are non-empty. Every perfect subset of a complete metric space is of size at least $|\mathbb{R}|$.

Again we interpolate between a large subset Y of a given metric space X and the real line using an ultrametric space. In this section we construct ultrametrics using some sort of infinite version of hierachically well-separated trees (see [1, Section 3.1]).

Definition 15 *We will work on the space 2^ω of all functions from ω to the set $2 = \{0, 1\}$. We approximate the elements of 2^ω using finite sequences of 0's and 1's, i.e., elements of the set $2^{<\omega} = \bigcup_{n=0}^{\infty} 2^n$. $2^{<\omega}$ is a tree with respect to the order "s is an initial segment of t". The points of 2^ω correspond to the infinite branches of the tree $2^{<\omega}$ via the map $f \mapsto \{f \restriction n : n \in \omega\}$.*

For $s \in 2^{<\omega}$ and $i \in 2$ let $s^\frown i$ denote the sequence obtained by extending s by the single digit i. For $f, g \in 2^\omega$ let $\mathrm{lci}(f, g)$ denote the longest common initial segment of f and g. We have $\mathrm{lci}(f, g) \in 2^{<\omega}$ if and only if f and g are distinct.

Lemma 16 *Let $K > 1$ and $\varepsilon = 1 - \frac{1}{K}$. Let $\Delta : 2^{<\omega} \to [0, \infty)$ be such that for* *all* $s \in 2^{<\omega}$ *we have*

$$\Delta(s^\frown 0), \Delta(s^\frown 1) \leq \frac{\varepsilon}{2} \cdot \Delta(s).$$

We define a metric on 2^ω *by letting* $d_\mathrm{u}(f,g) = \Delta(\mathrm{lci}(f,g))$.

Then d_u *is an ultrametric and* $(2^\omega, d_\mathrm{u})$ *is* K*-linear.*

PROOF. In order to verify that d_u is an ultrametric, let f, g and h be pairwise distinct elements of 2^ω. Let $s = \mathrm{lci}(f,g)$. If s is an initial segment of h, then s is also an initial segment of $\mathrm{lci}(f,h)$ and hence $d_\mathrm{u}(f,h) \leq \Delta(s) = d_\mathrm{u}(f,g)$. If s is not an initial segment of h, then $\mathrm{lci}(f,h) = \mathrm{lci}(g,h)$ and hence $d_\mathrm{u}(f,h) = \Delta(\mathrm{lci}(g,h)) = d_\mathrm{u}(g,h)$. In both cases we have

$$d_\mathrm{u}(f,h) \leq \max(d_\mathrm{u}(f,g), d_\mathrm{u}(g,h)),$$

showing that d_u indeed is an ultrametric.

We define an embedding of 2^ω into $\mathbb{R}$ by letting

$$e(f) = \sum_{n=0}^{\infty} (-1)^{f(n)} \cdot \frac{\Delta(f \restriction n)}{2}$$

for every $f \in 2^\omega$. The series $e(f)$ converges for every f since $(\Delta(f \restriction n))_{n\in\omega}$ decreases sufficiently fast. More precisely, for every $m \in \omega$,

$$\left| e(f) - \sum_{n=0}^{m} (-1)^{f(n)} \cdot \frac{\Delta(f \restriction n)}{2} \right| = \left| \sum_{n=m+1}^{\infty} (-1)^{f(n)} \cdot \frac{\Delta(f \restriction n)}{2} \right|$$
$$\leq \frac{\varepsilon}{2} \cdot \Delta(f \restriction m) \cdot \sum_{n=1}^{\infty} \frac{1}{2^n} = \frac{\varepsilon}{2} \cdot \Delta(f \restriction m).$$

It follows that if $f, g \in 2^\omega$ are distinct and $s = \mathrm{lci}(f,g)$ then

$$\frac{1}{K} \leq 1 - \varepsilon \leq \frac{(1-\varepsilon) \cdot d_\mathrm{u}(f,g)}{d_\mathrm{u}(f,g)}$$
$$= \frac{(1-\varepsilon) \cdot \Delta(s)}{d_\mathrm{u}(f,g)} \leq \frac{|e(f) - e(g)|}{d_\mathrm{u}(f,g)} \leq \frac{(1+\varepsilon) \cdot \Delta(s)}{d_\mathrm{u}(f,g)}$$
$$= \frac{(1+\varepsilon) \cdot d_\mathrm{u}(f,g)}{d_\mathrm{u}(f,g)} \leq 1 + \varepsilon \leq K.$$

Therefore e is a K-embedding. $\square$

Using Lemma 16 it is now easy to prove

Theorem 17 *Let $K > 1$. Let (X, d) be a complete metric space without isolated points. Then X has a perfect subset Y that K-embeds into an ultrametric space that is K-linear.*

PROOF. The proof of this theorem is a straight forward construction of a Cantor space using a tree of open sets.

Let $\varepsilon = 1 - \frac{1}{K}$. We choose a family $(x_s)_{s\in 2^{<\omega}}$ of points in X and a family $(O_s)_{s\in 2^{<\omega}}$ of open subsets of X such that the following conditions are satisfied:

(1) For all $s \in 2^{<\omega}$, $x_s \in O_s$.
(2) If $t \in 2^{<\omega}$ is a proper extension of $s \in 2^{<\omega}$, then $\mathrm{cl}(O_t) \subseteq O_s$.
(3) For all $s \in 2^{<\omega}$ the diameters of $U_{s^\frown 0}$ and $U_{s^\frown 1}$ are at most $\frac{\varepsilon}{4} \cdot \Delta(s)$ where $\Delta(s) = d(x_{s^\frown 0}, x_{s^\frown 1})$.

Since $\varepsilon < 1$, (3) implies

(4) If $s, t \in 2^{<\omega}$ are distinct sequences of the same length, then $\mathrm{cl}(U_s)$ and $\mathrm{cl}(U_t)$ are disjoint.

The families $(x_s)_{s\in 2^{<\omega}}$ and $(O_s)_{s\in 2^{<\omega}}$ can be chosen by recursion on the length of s since X has no isolated points and therefore every non-empty open subset of X is infinite.

By (1)–(3), for every $f : \omega \to 2$ the sequence $(x_{f\restriction n})_{n\in\omega}$ is Cauchy. Since X is complete, $x_f = \lim_{n\to\infty} x_{f\restriction n}$ exists. By (4), if $f \neq g$, then $x_f \neq x_g$. If follows that $e : 2^\omega \to X; f \mapsto x_f$ is 1-1. It is easily checked that $Y = e[2^\omega]$ is a perfect set. In fact, Y is a homeomorphic copy of the Cantor set.

Note that by (1)–(3), Δ satisfies the requirements of Lemma 16. Let d_{u} be the ultrametric on 2^ω defined from Δ. By Lemma 16, $(2^\omega, d_{\mathrm{u}})$ is K-linear.

It remains to show that e is a K-embedding with respect to d_{u} and d.

Let $f, g \in 2^\omega$ be distinct. Let $s = \mathrm{lci}(f, g)$. Then $d_{\mathrm{u}}(f, g) = \Delta(s)$. We may assume that $s^\frown 0$ is an initial segment of f and $s^\frown 1$ of g.

By (2), $x_f \in U_{s^\frown 0}$ and $x_g \in U_{s^\frown 1}$. Now by (3) we have

$$\frac{1}{K} \leq \frac{(1-\varepsilon)\cdot d(x_{s^\frown 0}, x_{s^\frown 1})}{d(x_{s^\frown 0}, x_{s^\frown 1})} \leq \frac{d(x_f, x_g)}{d_{\mathrm{u}}(x_f, x_g)} \leq \frac{(1+\varepsilon)\cdot d(x_{s^\frown 0}, x_{s^\frown 1})}{d(x_{s^\frown 0}, x_{s^\frown 1})} \leq K.$$

This shows that e indeed is a K-embedding. □

2.5 Compactness

We collect some properties of K-embeddability that are related to compactness. The results of this subsection will be used in Section 3 in order to analyze the infinite combinatorics of K-embeddability. A metric space X is *homogeneous* if for any two points $x, y \in X$ there is an isometry of X mapping x to y.

Theorem 18 *Let M be a separable metric space and let X be a homogeneous metric space in which every bounded set is contained in a compact set. If $K > 1$, then M K-embeds into X iff every finite subset of M K-embeds into X.*

PROOF. If M K-embeds into X, then so does every finite subset of M.

Now suppose that every finite subset of M K-embeds into X. Using the separability of M fix a dense subset $\{a_n : n \in \omega\}$ of M and set $F_n = \{a_0, \ldots, a_n\}$. For every n let e_n be a K-embedding of F_n into X. By the homogeneity of X we may assume that all e_n map a_0 to the same point x_0.

Note that for every $n \in \omega$, the sequence $(e_n(a_k))_{n \geq k}$ is bounded and therefore has a convergent subsequence. Inductively we can find an infinite subsequence $(e_{n_i})_{i \in \omega}$ of the sequence $(e_n)_{n \in \omega}$ such that for all $n \in \omega$ the sequence $(e_{n_i}(a_n))_{i \in \omega \wedge n_i \geq n}$ converges. For each $n \in \omega$ let

$$x_n = \lim_{i \to \infty} e_{n_i}(a_n).$$

It is easily checked that $a_n \mapsto x_n$ defines a K-embedding of $\{a_n : n \in \omega\}$ into X. This embedding has a unique continuous extension to all of M that is also a K-embedding. □

Corollary 19 *Let $K > 1$. If $n > m$, then there is a finite set $F \subseteq \mathbb{R}^n$ that is not K-embeddable into $\mathbb{R}^m$.*

PROOF. Since $\mathbb{R}^n$ is not homeomorphic to a subset of $\mathbb{R}^m$, $\mathbb{R}^n$ does not K-embed into $\mathbb{R}^m$. Now the existence of a finite set F that is not K-embeddable into $\mathbb{R}^m$ follows from Theorem 18. □

An argument similar to that in the proof of Theorem 18 yields the following:

Lemma 20 *Let M be a finite metric space and let X be a homogeneous metric space such that every bounded set is contained in a compact set. If M K-embeds into X for some K, then there is is a least such K.*

PROOF. Let $(K_n)_{n\in\omega}$ be a decreasing sequence of real numbers > 1 such that for all n, M K_n-embeds into X. Let $K = \lim_{n\to\infty} K_n$. For every $n \in \omega$ fix a K_n-embedding $e_n : M \to \mathbb{R}$. By the homogeneity of X we may assume that there is a point in M that is mapped to the same point in X by every e_n. Now we can thin out the sequence of embeddings e_n so that we obtain a subsequence $(e_{n_i})_{i\in\omega}$ with the property that for each $x \in M$ the sequence $(e_{n_i}(x))_{i\in\omega}$ converges. For each $x \in M$ let $e(x) = \lim_{i\to\infty} e_{n_i}(x)$. It is easily checked that $e : M \to \mathbb{R}$ is a K-embedding. □

Corollary 21 *Let M be a separable metric space and let X be a homogenous metric space such that every bounded subset of X is contained in a compact set. If for some K the space M K-embeds into X, then there is a least such K.*

PROOF. For each finite set $F \subseteq M$ let K_F denote the least K such that F K-embeds into X. Let

$$K = \sup\{K_F : F \subseteq M \text{ is finite}\}.$$

By Theorem 18, M K-embeds into X and K is minimal with this property. □

3 Covering numbers

3.1 Covering metric spaces by K-linear sets

Fix $K > 1$. Theorem 17 shows that every uncountable Polish space has large K-linear subset. We generalize this fact and show that it is consistent that every separable metric space can be covered by a small number of K-linear sets.

Let $\mathbb{U}$ be Urysohn's universal separable metric space. Since every separable metric space isometrically embeds into $\mathbb{U}$, it is sufficient to show that $\mathbb{U}$ can consistently be covered by a small number of K-linear sets. We will, however, carry out a forcing construction that works for any fixed separable metric space, not just the Urysohn space.

Definition 22 *Let M be a separable metric space. Let $\mathbb{P}_M$ denote the forcing notion consisting of finite subsets of M that are k-linear for some k with $1 < k < K$. The order on $\mathbb{P}_M$ is reverse inclusion.*

Lemma 23 *$\mathbb{P}_M$ is σ-linked.*

PROOF. First observe that all 2-element subsets of M isometrically embed into $\mathbb{R}$. It follows that the set of all singletons in $\mathbb{P}_M$ is linked.

Now let $n > 1$. For a partial 1-1 map e from M into $\mathbb{R}$ let Distortion(e) denote the least k such that the e is a k-embedding. If e is of size n, then e can be considered as an element of $(M \times \mathbb{R})^n$. Clearly, the map Distortion is continuous on the set of elements of $(M \times \mathbb{R})^n$ that correspond to 1-1 maps of size n.

Let $p \in \mathbb{P}_M$ be of size n, say $p = \{a_1, \dots, a_n\}$. Choose $k < K$ so that p is k-linear and fix a k-embedding $e : p \to \mathbb{R}$. Let $U_1, \dots, U_n \subseteq M$ be pairwise disjoint open sets such that for all $i \in \{1, \dots, n\}$, $a_i \in U_i$. Let $V_1, \dots, V_n \subseteq \mathbb{R}$ be pairwise disjoint open sets such that for all $i \in \{1, \dots, n\}$, $e(a_i) \in V_i$.

By the continuity of Distortion we can choose the U_i and V_i so small that for all $(x_1, \dots, x_n) \in U_1 \times \cdots \times U_n$ and all $(y_1, \dots, y_n) \in V_1 \times \cdots \times V_n$ we have

$$\text{Distortion}((x_1, y_1), \dots, (x_n, y_n)) < \frac{k+K}{2}.$$

We may assume that all the V_i are intervals of length ε for some fixed $\varepsilon > 0$. We may also assume that all the U_i are of diameter $< \varepsilon$. Finally, we may assume that the U_i are chosen from a fixed countable base of the topology on M.

Now, whenever

$$(x_1^0, \dots, x_n^0), (x_1^1, \dots, x_n^1) \in U_1 \times \cdots \times U_n$$

the conditions $\{x_1^0, \dots, x_n^0\}$ and $\{x_1^1, \dots, x_n^1\}$ are compatible: It is enough to show that $\{x_i^j : j \in 2 \wedge i \in \{1, \dots, n\}\}$ is $\frac{k+K}{2}$-linear. But we can construct a $\frac{k+K}{2}$-embedding

$$f : \{x_i^j : j \in 2 \wedge i \in \{1, \dots, n\}\} \to \mathbb{R}$$

as follows:

For $i \in \{1, \dots, n\}$ we choose y_i^0 and y_i^1 in V_i of the same distance as x_i^0 and x_i^1. This is possible since the interval V_i is of length ε. We claim that the map f that maps each x_i^j to y_i^j is a $\frac{k+K}{2}$-embedding.

To see this, consider two distinct points $a, b \in \text{dom}(f)$. If $a \in U_{i_a}$ and $b \in U_{i_b}$ with $i_a \neq i_b$, then $f(a) \in V_{i_a}$ and $f(b) \in V_{i_b}$. From the choice of the U_i and the V_i it follows that

$$\frac{2}{k+K} \leq \frac{d_M(f(a), f(b))}{|a-b|} \leq \frac{k+K}{2}.$$

If a and b lie in the same U_i, then

$$\frac{d_M(f(a), f(b))}{|a-b|} = 1.$$

It follows that f indeed is a $\frac{k+K}{2}$-embedding.

This argument shows that for each condition $p \in \mathbb{P}_M$ with $p = \{x_1, \ldots, x_n\}$ for some pairwise distinct x_i, there are pairwise disjoint basic open sets $U_1, \ldots, U_n \subseteq M$ such that the set $P_{U_1,\ldots,U_n}$ of all conditions $q = \{y_1, \ldots, y_n\} \in \mathbb{P}_M$ with $(y_1, \ldots, y_n) \in U_1 \times \cdots \times U_n$ is linked and $p \in P_{U_1,\ldots,U_n}$. Since there are only countably many basic open sets and hence only countably many finite sequences of those, it follows that $\mathbb{P}_M$ is σ-linked. □

Using this lemma it is easy to show

Theorem 24 *There is a c.c.c. forcing notion $\mathbb{P}$ such that*

$$\Vdash_{\mathbb{P}} \text{“for every } K > 1 \text{ and every separable metric space } M, \quad M \text{ is coverable by at most } \aleph_1 \text{ } K\text{-linear sets”}$$

PROOF. Since every separable metric space is isometric to a subspace of Urysohn's universal space $\mathbb{U}$, it is enough to force that $\mathbb{U}$ can be covered by at most $\aleph_1$ K-linear sets. It is worth pointing out that the space $\mathbb{U}$ in a forcing extension $V[G]$ of the set-theoretic universe V is simply the completion in $V[G]$ of the space $\mathbb{U}$ in V.

We now construct $\mathbb{P}$ as follows:

Let $\mathbb{Q}$ denote the finite support product of countably many copies of $\mathbb{P}_{\mathbb{U}}$. $\mathbb{Q}$ is σ-linked:

Let $\mathbb{P}_{\mathbb{U}} = \bigcup_{n\in\omega} L_n$ where each L_n is linked. For each $f \in \omega^{<\omega}$ let L_f denote the subset of $\mathbb{Q}$ that consists of all conditions q such that $q(n) = L_{f(n)}$ whenever $n \in \text{dom}(f)$ and $q(n) = 1_{\mathbb{P}_{\mathbb{U}}}$ otherwise. It is easily checked that each L_f is linked and that $\mathbb{Q} = \bigcup_{f\in\omega^{<\omega}} L_f$. It follows that $\mathbb{Q}$ is σ-linked.

$\mathbb{Q}$ generically adds countably many K-linear subsets of $\mathbb{U}$ and an easy density argument shows that these sets cover the ground model version of $\mathbb{U}$. Now let $\mathbb{P}$ be the finite support iteration of $\mathbb{Q}$ of length ω_1.

Let G be a $\mathbb{P}$-generic filter over the ground model V. Since the length of the iteration is of uncountable cofinality, every new real is added at some intermediate stage of the iteration. In particular, every element of the Urysohn space of the final extension appears at some intermediate stage of the iteration.

It follows that the $\aleph_1$ K-linear subsets of $\mathbb{U}$ that have been added in the process, countably many at each stage, cover the whole Urysohn space of $V[G]$ $\square$

Corollary 25 *It is consistent with arbitrarily large values of $2^{\aleph_0}$ that every separable metric space can be covered by $\aleph_1$ K-linear sets.*

PROOF. Start with a model of set theory with the desired size of $2^{\aleph_0}$ and force with the partial order $\mathbb{P}$ constructed in Theorem 24. Since $\mathbb{P}$ is c.c.c., no cardinals are collapsed. Since $\mathbb{P}$ has a dense subset of size $2^{\aleph_0}$, the value of $2^{\aleph_0}$ is not changed in the generic extension by $\mathbb{P}$. $\square$

Note that the argument can be easily modified in such a way that we obtain the consistency result for all $K > 1$ at the same time. All that has to be guaranteed is that for every $n > 0$, the constant $K = 1 + \frac{1}{n}$ is dealt with cofinally often during the iteration.

It is worth pointing out that even $\mathbb{R}^2$ is not coverable by less than $2^{\aleph_0}$ isometric copies of subsets of $\mathbb{R}$. This is because every isometric copy of a subset of $\mathbb{R}$ is contained in a one-dimensional affine subspace of $\mathbb{R}^2$ and less than $2^{\aleph_0}$ one-dimensional affine subspaces do not cover $\mathbb{R}^2$.

3.2 Covering $\mathbb{R}^n$ by low-distortion copies of the real line

Theorem 24 might be regarded as slightly unsatisfactory since the K-linear sets added by forcing notions of the form $\mathbb{P}_M$ are very thin, i.e., they actually do not resemble anything that really looks like the real line.

It may be more natural to try to cover separable metric spaces by bi-Lipschitz-images of $\mathbb{R}$. There are obvious limitations to this: in a zero-dimensional space such as the Cantor space every continuous image of the real line is just a singleton. Hence we cannot hope to cover general separable metric spaces by less than $2^{\aleph_0}$ continuous images of the real line.

On the other hand, there is no obvious limitation if we want to cover $\mathbb{R}^n$ by a small number of bi-Lipschitz images of $\mathbb{R}$. The problem is, however, that K-bi-Lipschitz maps from closed subsets of $\mathbb{R}$ into $\mathbb{R}^n$ might not be extendable to K-bi-Lipschitz maps that are defined on all of $\mathbb{R}$.

We will produce K-bi-Lipschitz images of $\mathbb{R}$ using graphs of Lipschitz functions.

Lemma 26 *Let $f : \mathbb{R} \to \mathbb{R}$ be Lipschitz of constant $\varepsilon > 0$. Then f, regarded as a subset of $\mathbb{R}^2$, is a $\sqrt{1+\varepsilon}$-bi-Lipschitz image of $\mathbb{R}$.*

PROOF. Let $g : \mathbb{R} \to \mathbb{R}^2$ be defined by letting $g(x) = (f(x), x)$. Now, for two distinct $x, y \in \mathbb{R}$ we have

$$\frac{|g(x) - g(y)|}{|x-y|} = \frac{\sqrt{(f(x) - f(y))^2 + (x-y)^2}}{|x-y|} \geq 1.$$

On the other hand,

$$\frac{|g(x) - g(y)|}{|x-y|} = \frac{\sqrt{(f(x) - f(y))^2 + (x-y)^2}}{|x-y|}$$
$$\leq \frac{\sqrt{(\varepsilon \cdot (x-y))^2 + (x-y)^2}}{|x-y|} = \sqrt{1+\varepsilon}$$

Hence g is $\sqrt{1+\varepsilon}$-bi-Lipschitz. Clearly, $f = g[\mathbb{R}]$. □

We use the Dual Open Coloring Axiom (DOCA) to show the consistency of covering $\mathbb{R}^n$ by fewer than $2^{\aleph_0}$ K-bi-Lipschitz images of $\mathbb{R}$.

For every topological space X the set $[X]^2$ of two element subsets of X carries the topology generated by sets of the form

$$\{\{x, y\} : x \in U \wedge y \in V\}$$

where U and V are disjoint open subsets of X. An *open pair cover on X* is a finite collection $U_1, \dots, U_n$ of open subsets of $[X]^2$ such that $[X]^2 = U_1 \cup \dots \cup U_n$. A set $H \subseteq X$ is *homogeneous of color $i \in \{1, \dots, n\}$* if $[H]^2 \subseteq U_i$. Now DOCA is the statement

"For every open pair cover $U_1, \dots, U_n$ on a Polish space X, X is the union of fewer than $2^{\aleph_0}$ homogeneous sets".

The axiom DOCA is known to be consistent with ZFC [4] and it implies $2^{\aleph_0} > \aleph_1$. All known models of DOCA satisfy $2^{\aleph_0} = \aleph_2$.

Theorem 27 *For every $K > 1$ and every $n > 1$, DOCA implies that $\mathbb{R}^2$ is the union of less than $2^{\aleph_0}$ K-bi-Lipschitz images of $\mathbb{R}$.*

PROOF. We define an open pair cover on $\mathbb{R}^2$ as follows:

Let $\varepsilon > 0$ be such that $\sqrt{1+\varepsilon} < K$. Let $\gamma = \arctan \varepsilon$. Choose finitely many open intervals $I_1, \dots, I_n$ of angles such that every angle is contained in at least one I_i and moreover, each I_i has diameter less than 2γ. For each $i \in \{1, \dots, n\}$ let U_i denote the set of all

$$\{(x_1, y_1), (x_2, y_2)\} \in [\mathbb{R}^2]^2$$

such that the angle of the line through (x_1, y_1) and (x_2, y_2) with the x-axis is an element of I_i.

Since the I_i are open, the U_i are open subsets of $[\mathbb{R}^2]^2$. For each i let α_i be the midpoint of the interval I_i. Now, if H is a homogeneous subset of $\mathbb{R}^2$ of color i, then rotating H by the angle $-\alpha_i$ yields a subset of $\mathbb{R}^2$ that is the graph of a partial Lipschitz map from $\mathbb{R}$ to $\mathbb{R}$ of constant ε.

A partial Lipschitz map of constant ε has a unique extension to the closure of its domain that is also Lipschitz of constant ε. A partial Lipschitz map of constant ε from a closed subset of $\mathbb{R}$ to $\mathbb{R}$ can be extended to all of $\mathbb{R}$ by linear interpolation on the open intervals where the map is originally undefined. By Lemma 26 and by the choice of K, the graph of a Lipschitz map of constant ε from $\mathbb{R}$ to $\mathbb{R}$ is a K-bi-Lipschitz image of $\mathbb{R}$.

Being a K-bi-Lipschitz image of $\mathbb{R}$ is clearly rotation invariant. It follows that every homogeneous set H for the open pair cover $U_1, \dots, U_n$ is contained in a K-bi-Lipschitz image of $\mathbb{R}$. Using the Dual Open Coloring Axiom we obtain a family $\mathcal{H} \subseteq \mathcal{P}(\mathbb{R}^2)$ of size $< 2^{\aleph_0}$ consisting of homogeneous sets such that $\mathbb{R}^2 = \bigcup \mathcal{H}$. Since each $H \in \mathcal{H}$ is contained in a K-bi-Lipschitz image of $\mathbb{R}$ by the argument above, $\mathbb{R}^2$ is coverable by $< 2^{\aleph_0}$ K-bi-Lipschitz images of $\mathbb{R}$. $\square$

The next lemma shows that higher dimensional instances of Theorem 27 actually follow from it.

Lemma 28 *Let $K > 1$. If $\mathbb{R}^2$ can be covered by κ K-linear sets for some $K > 1$ and some cardinal κ, then for every $n > 1$, $\mathbb{R}^n$ can be covered by κ K^{n-1}-linear sets.*

If $\mathbb{R}^2$ can be covered by κ K-bi-Lipschitz images of $\mathbb{R}$ for some $K > 1$ and some cardinal κ, then for every $n > 1$, $\mathbb{R}^n$ can be covered by κ K^{n-1}-bi-Lipschitz images of $\mathbb{R}$.

PROOF. We show the lemma by induction on n and concentrate on the statement about K-bi-Lipschitz images of $\mathbb{R}$. The K-linear sets can be handled in the same way.

Suppose $\mathbb{R}^n$ can be covered by less than $2^{\aleph_0}$ K^{n-1}-bi-Lipschitz images of $\mathbb{R}$. Fix a family $\mathcal{F}$ of size $< 2^{\aleph_0}$ of K^{n-1}-bi-Lipschitz maps from $\mathbb{R}$ to $\mathbb{R}^n$ whose images cover all of $\mathbb{R}^n$. For each $f \in \mathbb{F}$ the map

$$\mathrm{id} \times f : \mathbb{R}^2 \to \mathbb{R} \times \mathbb{R}^n; (x, y) \mapsto (x, f(y))$$

is clearly K^{n-1}-bi-Lipschitz.

Fix a family $\mathcal{G}$ of size $< 2^{\aleph_0}$ of K-bi-Lipschitz maps from $\mathbb{R}$ to $\mathbb{R}^2$ whose images cover all of $\mathbb{R}^2$. For all $g \in \mathcal{G}$ and all $f \in \mathcal{F}$ the map

$$(\mathrm{id} \times f) \circ g : \mathbb{R} \to \mathbb{R}^{n+1}$$

is K^n-bi-Lipschitz. Moreover, the images of all the maps $(\mathrm{id} \times f) \circ g$, $f \in \mathcal{F}$, $g \in \mathcal{G}$, cover all of $\mathbb{R}^{n+1}$. But there are only $|\mathcal{F}| \cdot |\mathcal{G}| < 2^{\aleph_0}$ maps of this form. □

Corollary 29 *For every $n > 1$ and every $K > 1$, DOCA implies that $\mathbb{R}^n$ can be covered by less than $2^{\aleph_0}$ K-bi-Lipschitz copies of $\mathbb{R}$.*

A more general form of Lemma 28 can be proved in the same way.

Lemma 30 *Let $K > 1$, $n > m \geq 1$ and let κ be a cardinal. If $\mathbb{R}^{m+1}$ can be covered by κ sets that K-embed into $\mathbb{R}^m$, then $\mathbb{R}^n$ can be covered by κ sets that K^{n-m}-embed into $\mathbb{R}^m$.*

If $\mathbb{R}^{m+1}$ can be covered by κ K-bi-Lipschitz images of $\mathbb{R}^m$, then $\mathbb{R}^n$ can be covered by κ K^{n-m}-bi-Lipschitz images of $\mathbb{R}^m$.

4 Covering the plane by K-linear sets and localization numbers

We use the results from Section 2.5 to show a relation between the so-called localization numbers and coverings of $\mathbb{R}^2$ by K-linear sets.

Definition 31 *Let $m > n > 0$. A set $S \subseteq m^\omega$ is* n-ary *if no $n+1$ distinct points of S pairwise disagree for the first time at the same coordinate. The least number of n-ary sets that cover m^ω is the* localization number $\mathfrak{l}_{n,m}$.

It is not hard to see that

$$2^{\aleph_0} = \mathfrak{l}_{1,2} \geq \mathfrak{l}_{2,3} \geq \dots$$

Moreover, a simple induction shows that $\mathfrak{l}_{n,m} = \mathfrak{l}_{n,n+1}$ whenever $m > n > 0$. For all $n, m > 1$ with $n < m$ the statement $\mathfrak{l}_{m,m+1} > \mathfrak{l}_{n,n+1}$ is consistent with

ZFC [14]. It is even possible to separate finitely many numbers $\mathfrak{l}_{n,n+1}$ in the same model of set theory [6].

Theorem 32 *Let $K > k > 1$, $n > 1$ and suppose there is an $(n+1)$-element subset F of $\mathbb{R}^2$ that is not K-linear. Then at least $\mathfrak{l}_{n,n+1}$ k-linear sets are necessary to cover $\mathbb{R}^2$.*

PROOF. Let $F = \{x_0, \ldots, x_n\}$. There are pairwise disjoint open sets

$$O_0, \ldots, O_n \subseteq \mathbb{R}^2$$

such that $x_i \in O_i$ for all $i \leq n$ and whenever $y_i \in O_i$ for all $i \leq n$, then $\{y_0, \ldots, y_n\}$ is k-linear. Using F, respectively the collection $U_0, \ldots, U_n$, as a template, we construct an embedding $e : (n+1)^\omega \to \mathbb{R}^2$ such that whenever $S \subseteq (n+1)^\omega$ is not n-ary, then $e[S]$ is not k-linear.

It is clear that the existence of e gives the desired result: If $\mathcal{S}$ is a family of k-linear subsets of $\mathbb{R}^2$ such that $\bigcup \mathcal{S} = \mathbb{R}^2$, then $\{e^{-1}[S] : S \in \mathcal{S}\}$ is a family of n-ary sets that covers $(n+1)^\omega$.

Let G denote the set $\mathrm{cl}(U_0 \cup \cdots \cup U_n)$. By recursion on $(n+1)^{<\omega}$ we define a family $(f_\sigma)_{\sigma \in (n+1)^{<\omega}}$ of affine linear bijections on $\mathbb{R}^2$. Each f_σ will be the composition of a translation and dilation, i.e., a homothetic transformation. Let $f_\emptyset$ be the identity on $\mathbb{R}^2$. Suppose f_σ has been defined for some $\sigma \in (n+1)^{<\omega}$. For each $i \leq n$ choose a homothetic transformation $f_{\sigma^\frown i} : \mathbb{R}^2 \to \mathbb{R}^2$ that maps the set G to a subset of $f_\sigma[U_i]$. Moreover, choose $f_{\sigma^\frown i}$ so that the diameter of each $f_{\sigma^\frown i}[U_j]$, $j \leq n$, is at most one half of the diameter of $f_\sigma[U_i]$.

Now we are ready to define the embedding e. For $a \in (n+1)^\omega$ let

$$e(a) = \lim_{k \to \infty} f_{a \restriction k}(x_{a(k)}).$$

Equivalently, $e(a)$ is the unique element of the set $\bigcap_{k \in \omega} f_{a \restriction k}[\mathrm{cl}(U_{a(k)})]$.

If $S \subseteq (n+1)^\omega$ is not n-ary, then there are points $a_0, \ldots, a_n \in S$ that pairwise disagree for the first time on the same coordinate k, i.e., for some $\sigma \in (n+1)^k$, $\sigma = a_0 \restriction k = \cdots = a_n \restriction k$ and $a_0(k), \ldots, a_n(k)$ are pairwise different. We may assume that $a_i(k) = i$ for all $i \leq n$. Now

$$e(a_i) \in f_\sigma[U_{a_i(k)}] = f_\sigma[U_i].$$

Since f_σ is a homothetic transformation, by the choice of the original set F and by the choice of the U_i, the set $\{e(a_0), \ldots, e(a_n)\}$ is not k-linear. This finishes the proof of the Theorem. $\square$

A similar argument shows that for all $n \geq 2$ and all sufficiently small $K > 1$, at least $\mathfrak{l}_{n,n+1}$ subsets of $\mathbb{R}^n$ that K-embed into $\mathbb{R}^{n-1}$ are necessary to cover all of $\mathbb{R}^n$.

The next theorem nicely complements Theorem 32.

Theorem 33 *Let $K > k > 1$ and $n > 1$ be such that every $(n+1)$-element subset of $\mathbb{R}^2$ is k-linear. Then there is a forcing extension of the set-theoretic universe in which $\mathbb{R}^2$ can be covered by $\aleph_1$ K-linear subsets while $\mathfrak{l}_{n,n+1} > \aleph_1$.*

We first need to make sure that for every $n > 1$ there is some $k > 1$ such that every $(n+1)$-element subset of $\mathbb{R}^2$ is k-linear. This follows from

Theorem 34 (Matoušek [13]) *For every $n > 0$ there is $K_n > 1$ such that every metric space with n points is K_n-linear.*

There are two different strategies to prove Theorem 33. The first one is to start with a model of set theory in which CH holds and then to increase $\mathfrak{l}_{n,n+1}$ in a way that tends to increase other cardinal invariants as little as possible. This approach is resonably well understood [15].

We will use a different approach. Namely, we start with a model of set theory in which $\mathfrak{l}_{n,n+1}$ is large, for example a model of Martin's Axiom and $\neg$CH, and then decrease the number of K-linear sets needed to cover $\mathbb{R}^2$ without decreasing $\mathfrak{l}_{n,n+1}$.

A property of a forcing notion $\mathbb{P}$ that guarantees that $\mathfrak{l}_{n,n+1}$ is not decreased by forcing with $\mathbb{P}$ is σ-$(n+1)$-linkness.

Definition 35 *A subset S of a forcing notion $\mathbb{P}$ is n-*linked *if any n elements of S have a common extension in $\mathbb{P}$. $\mathbb{P}$ is σ-n-*linked *if $\mathbb{P}$ is the union of countably many n-linked subsets.*

Lemma 36 *Suppose $\mathbb{P}$ is σ-$(n+1)$-linked. Let G be $\mathbb{P}$-generic over the ground model V. Then $(\mathfrak{l}_{n,n+1})^{V[G]} \geq (\mathfrak{l}_{n,n+1})^V$.*

PROOF. First observe that $\mathbb{P}$ is c.c.c. Hence $V[G]$ has the same cardinals as V. Now let κ be a cardinal and assume that in $V[G]$, $(S_\alpha)_{\alpha<\kappa}$ is a family of n-ary subsets of $(n+1)^\omega$ with $(n+1)^\omega = \bigcup_{\alpha\in\kappa} S_\alpha$. For each $\alpha < \kappa$ let $\dot{S}_\alpha$ be a $\mathbb{P}$-name for S_α.

There is a condition in G that forces $(n+1)^\omega$ to be the union of the $\dot{S}_\alpha$, $\alpha < \kappa$, and that forces each $\dot{S}_\alpha$ to be n-ary. For simplicity we assume that the largest element of $\mathbb{P}$ already forces this. The general case is handled by exactly the same argument but is notationally slightly more complicated.

We work in the ground model. Let $\mathbb{P} = \bigcup_{i\in\omega} P_i$ with each P_i $(n+1)$-linked. For every $i \in \omega$ and every $\alpha < \kappa$ let

$$S^i_\alpha = \{a \in (n+1)^\omega : \exists p \in P_i(p \Vdash a \in \dot{S}_\alpha)\}.$$

For each $a \in (n+1)^\omega$ there is some $\alpha < \kappa$ and a condition $p \in \mathbb{P}$ that forces a to be an element of $\dot{S}_\alpha$. If $p \in P_i$, then $a \in S^i_\alpha$. It follows that $(n+1)^\omega = \bigcup_{\alpha<\kappa\wedge i\in\omega} S^i_\alpha$. The proof of the lemma is finished if we can show that each S^i_α is n-ary. This is because in V, $(n+1)^\omega$ is the union of the κ n-ary sets S^i_α. This shows $(\mathfrak{l}_{n,n+1})^V \leq (\mathfrak{l}_{n,n+1})^{V[G]}$.

Now let $\alpha < \kappa$ and $i \in \omega$. Suppose $a_0, \dots, a_n$ are elements of S^i_α. Fix $p_0, \dots p_n \in P_i$ such that for all $j \leq n$, p_j forces a_j to be in $\dot{S}_\alpha$. Since P_i is $(n+1)$-linked, the p_j have a common extension $p \in \mathbb{P}$. Now p forces $a_0, \dots, a_n$ to be elements of $\dot{S}_\alpha$. Since $\dot{S}_\alpha$ is forced to be n-ary, $a_0, \dots, a_n$ cannot pairwise disagree for the first time at the same coordinate. Hence S^i_α is indeed n-ary. $\square$

Lemma 37 *Suppose K, k and n are as in Theorem 33. Then there is a forcing notion $\mathbb{P}$ that is σ-$(n+1)$-linked and adds countably many k-linear subsets of $\mathbb{R}^2$ that cover the ground model plane.*

PROOF. As in the proof of Theorem 24, the basic building block is the forcing notion $\mathbb{P}_{\mathbb{R}^2}$ as defined in Definition 22. $\mathbb{P}_{\mathbb{R}^2}$ is the set of finite subsets of $\mathbb{R}^2$ that are ℓ-linear for some $\ell < K$, ordered by reverse inclusion. Lemma 23 states that $\mathbb{P}_{\mathbb{R}^2}$ is σ-linked.

The only modifaction that has to be applied to the proof of Lemma 23 to give σ-$(n+1)$-linkedness in the current situation is this:

Suppose $\ell < K$. Given a finite ℓ-linear subset $\{x_1, \dots, x_m\}$ of $\mathbb{R}^2$, there are pairwise disjoint open neighborhoods U_j of the x_j such that if for all $j \in \{1, \dots, m\}$, F_j is an $(n+1)$-element subset of U_j, then $\bigcup_{1\leq j\leq m} F_j$ is ℓ'-linear for some ℓ' with $k < \ell' < K$. This follows by choosing the U_j sufficiently small and using the fact that each F_j is k-linear, where $k < K$.

This implies that whenever $p_i = \{y^i_1, \dots, y^i_m\}$, $i \leq n$, are conditions in $\mathbb{P}_{\mathbb{R}^2}$ with $y^i_j \in U_j$ for all $i \leq n$ and $1 \leq j \leq m$, then the p_i have a common extension in $\mathbb{P}_{\mathbb{R}^2}$. Now the same argument as in the proof of Lemma 23 shows that $\mathbb{P}_{\mathbb{R}^2}$ is σ-$(n+1)$-linked.

We now define $\mathbb{P}$ to be the finite support product of countably many copies of $\mathbb{P}_{\mathbb{R}^2}$. An easy density argument shows that $\mathbb{P}$ indeed adds countably many K-linear sets that cover the ground model plane. As with the corresponding statement for σ-linkedness in the proof of Theorem 24, $\mathbb{P}$ is σ-$(n{+}1)$-linked. $\square$

Lemma 38 *An iteration of σ-n-linked forcing notions of length $< (2^{\aleph_0})^+$ is σ-n-linked.*

PROOF. Let δ be an ordinal below $(2^{\aleph_0})^+$ and suppose $((\mathbb{P}_\alpha)_{\alpha\leq\delta}, (\dot{\mathbb{Q}}_\alpha)_{\alpha<\delta})$ is an iteration of σ-n-linked forcing notions, i.e., assume that for each $\alpha < \delta$,

$$\Vdash_{\mathbb{P}_\alpha} \dot{\mathbb{Q}}_\alpha \text{ is } \sigma\text{-}n\text{-linked.}$$

For each $\alpha < \delta$ fix $\mathbb{P}_\alpha$-names $\dot{Q}^i_\alpha$, $i \in \omega$, such that

$$\Vdash_{\mathbb{P}_\alpha} \dot{\mathbb{Q}}_\alpha = \bigcup_{i\in\omega} \dot{Q}^i_\alpha.$$

Clearly, it is enough to show that $\mathbb{P}_\delta$ has a dense subset that is σ-n-linked. By induction on the length of the iteration it is easy to show that $\mathbb{P}_\delta$ has a dense subset P consisting of conditions p such that for all $\alpha < \delta$, if α is in the support of p, then $p \restriction \alpha$ decides which $\dot{Q}^i_\alpha$ contains p.

By the Hewitt-Marczewski-Pondiczery Theorem (see [3, Theorem 2.3.15]) the space ω^δ is separable and hence there is a countable family $\mathcal{D}$ of functions from δ to ω such that every finite partial function from δ to ω has an extension in $\mathcal{D}$.

Now for each $f \in \mathbb{D}$ we define a subset P_f of P as follows: a condition $p \in P$ belongs to P_f if for all $\alpha < \delta$,

$$p \restriction \alpha \Vdash p(\alpha) \in \dot{Q}^{f(\alpha)}_\alpha.$$

By the choice of $\mathcal{D}$ and P, $P = \bigcup_{f\in\mathcal{D}} P_f$. We are finished if we can show that each P_f is n-linked.

Fix $f \in \mathcal{D}$ and let $p_1, \dots, p_n \in P_f$. By recursion on $\alpha < \delta$ we define a condition p such that for all $\alpha < \delta$,

$$p \restriction \alpha \Vdash p(\alpha) < p_1(\alpha), \dots, p_n(\alpha).$$

At stage α, a suitable name $p(\alpha)$ exists since $p \restriction \alpha$ is a common extension of $p_1 \restriction \alpha, \dots, p_n \restriction \alpha$ and thus

$$p \restriction \alpha \Vdash p_1(\alpha), \dots, p_n(\alpha) \in \dot{Q}^{f(\alpha)}_\alpha.$$

□

Proof of Theorem 33 We start by forcing $\mathfrak{l}_{n,n+1} > \aleph_1$. Since each n-ary subset of $(n+1)^\omega$ is nowhere dense, $\mathfrak{l}_{n,n+1} > \aleph_1$ follows for instance from

Martin's Axiom for $\aleph_1$ dense sets. Now we perform an iteration of length $\aleph_1$ of a forcing notion $\mathbb{P}$ as in Lemma 37. By Lemma 38, this iteration is still σ-$(n+1)$-linked and hence, by Lemma 36, the $\mathfrak{l}_{n,n+1}$ of the final model is still $> \aleph_1$.

On the other hand, during the final iteration $\aleph_1$ k-linear sets have been added, countably many at each stage, that together cover the plane of the final model of set theory. $\square$

A similar proof shows that it is consistent that for every $K > 1$, $\mathbb{R}^{n+1}$ can be covered by less than $\mathfrak{l}_{n,n+1}$ subsets that K-embed into $\mathbb{R}^n$.

Corollary 39 *There is a strictly increasing sequence $(K_n)_{n\in\omega}$ of real numbers > 1 such that whenever $n < m$, then there is a forcing extension of the set-theoretic universe where $\mathbb{R}^2$ can be covered by $\aleph_1$ K_m-linear set, but not by $\aleph_1$ K_n-linear sets.*

PROOF. For each $n \geq 3$ let

$$C_n = \{c > 1 : \text{Every } n\text{-element subset of } \mathbb{R}^2 \text{ is } c\text{-linear}\}.$$

By Theorem 34, C_n is non-empty. By Lemma 20, C_n is the intersection of a family of closed sets and hence has a minimal element c_n. By Corollary 19, the set $C = \{c_n : n \geq 3\}$ is unbounded in $\mathbb{R}$. Let $(K_n)_{n\in\omega}$ be a strictly increasing sequence of real numbers > 1 such that for all $n \in \omega$ there is some $c \in C$ such that $K_n < c < K_{n+1}$.

Now, if $n < m$, then there is some i such that if c minimal with the property that every $(i+1)$-element subset of $\mathbb{R}^2$ is c-linear, then $K_n < c < K_m$. By Theorem 33 there is a forcing extension of the universe in which $\mathbb{R}^2$ can be covered by $\aleph_1$ K_m-linear sets while $\mathfrak{l}_{i,i+1} > \aleph_1$. By Lemma 32, at least $\mathfrak{l}_{i,i+1}$ K_n-linear sets are neccessary to cover $\mathbb{R}^2$. $\square$

4.1 Discussion

For each $K > 1$ let $\kappa(K)$ denote the least number of K-linear subsets of $\mathbb{R}^2$ that cover $\mathbb{R}^2$. Using the models constructed in [6] it is actually possible to separate finitely many of the cardinal invariants $\kappa(K_n)$ at the same time. It is totally open is whether for every two real numbers k, K with $1 < k < K$ there is a forcing extension of the universe in which $\kappa(K) < \kappa(k)$. If yes, this would give a family of cardinal invariants of order type $\mathbb{R}$, where one cardinal invariant κ is considered to be strictly below another cardinal invariant λ if

$\kappa \leq \lambda$ holds in every generic extension of the set-theoretic universe and there is a generic extension of the universe where $\kappa < \lambda$. A negative answer would probably come with some interesting geometrical insight.

References

[1] Y. Bartal, N. Linial, M. Mendel, A. Naor, *On Metric Ramsey Type Phenomena*, Annals of Mathematics 162 (2005), 643–709

[2] J. Bourgain, T. Figiel, V. Milman, *On Hilbertian subsets of finite metric spaces*, Israel J. Math., 55 (1986), 147–152

[3] R. Engelking, *General topology*, translated from the Polish by the author, Second edition, Sigma Series in Pure Mathematics, Heldermann Verlag, Berlin (1989)

[4] S. Geschke, *A dual open coloring axiom*, Annals of Pure and Applied Logic 140 (2006), 40–51

[5] S. Geschke, M. Goldstern, M. Kojman, *Continuous Ramsey theory on Polish spaces and covering the plane by functions,* Journal of Mathematical Logic, Vol. 4 (2004), No. 2, 109–145

[6] S. Geschke, M. Kojman, *Convexity numbers of closed subsets in* $\mathbb{R}^n$, Proc. Am. Math. Soc. 130 (2002), No.10, 2871–2881

[7] R. Graham, B. Rothschild, J. Spencer, *Ramsey theory*, John Wiley & Sons Inc., New York, second edition (1990)

[8] T. Jech, *Set theory*, the third millennium edition, revised and expanded, Springer Monographs in Mathematics, Springer Berlin (2003)

[9] H. Karloff, Y. Rabani, Y. Ravid, *Lower bounds for randomized k-server and motion planning algorithms*, SIAM J. Comput., 23 (1994), 293–312

[10] A. Kechris, *Classical Descriptive Set Theory*, Graduate Texts in Mathematics, Springer (1995)

[11] K. Kunen, *Set theory. An introduction to independence proofs*, 2nd print, Studies in Logic and the Foundations of Mathematics, 102, Amsterdam-New York-Oxford: North-Holland (1983)

[12] J. Matoušek, *Ramsey-like properties for bi-Lipschitz mappings of finite metric spaces*, Commentat. Math. Univ. Carol. 33, No.3 (1992), 451–463

[13] J. Matoušek. *Bi-lipschitz embeddings into low-dimensional euclidean spaces*, Comment. Math. Univ. Carolinae 31 (1990), 589–600

[14] L. Newelski, A. Rosłanowski, *The ideal determined by the unsymmetric game* Proc. Amer. Math. Soc. 117 (1993), no. 3, 823–831

[15] J. Zapletal, *Descriptive Set Theory and Definable Forcing*, Memoirs of the Am. Math. Soc. 793 (2004)