Topology and its Applications

ELSEVIER

Cosmic dimensions

Alan Dow [a,1], Klaas Pieter Hart [b,*]

[a] *Department of Mathematics, UNC – Charlotte, 9201 University City Blvd., Charlotte, NC 28223-0001, USA*
[b] *Faculty of Electrical Engineering, Mathematics, and Computer Science, TU Delft, Postbus 5031, 2600 GA Delft, The Netherlands*

Abstract

Martin's Axiom for σ-centered partial orders implies that there is a cosmic space with non-coinciding dimensions.
© 2007 Elsevier B.V. All rights reserved.

MSC: primary 54F45; secondary 03E50, 54E20

Keywords: Cosmic space; Covering dimension; Small inductive dimension; Large inductive dimension; Martin's Axiom

Introduction

A fundamental result in dimension theory states that the three basic dimension functions, dim, ind and Ind, coincide on the class of separable metrizable spaces. Examples abound to show that this does not hold in general outside this class. In [1] Arkhangel′skiĭ asked whether the dimension functions coincide on the class of cosmic spaces. These are the regular **co**ntinuous images of **s**eparable **m**etric spaces and they are characterized by the conjunction of regularity and having a countable network, see [5]. A *network* for a topological space is a collection of (arbitrary) subsets such that every open set is the union of some subfamily of that collection. In [7] Vedenisoff proved that ind and Ind coincide on the class of perfectly normal Lindelöf spaces, see also [3, Section 2.4]. As the cosmic spaces belong to this class Arkhangel′skiĭ's question boils down to whether dim = ind for cosmic spaces.

In [2] Delistathis and Watson constructed, assuming the Continuum Hypothesis, a cosmic space X with $\dim X = 1$ and $\operatorname{ind} X \geqslant 2$; this gave a consistent negative answer to Arkhangel′skiĭ's question.

The purpose of this paper is to show that the example can also be constructed under the assumption of Martin's Axiom for σ-centered partial orders. The overall strategy is that of [2]: we refine the Euclidean topology of a one-dimensional subset X of the plane to get a topology τ with a countable network, such that $\dim(X, \tau) = 1$ and in which the boundary of every non-dense open set is (at least) one-dimensional, so that $\operatorname{ind}(X, \tau) \geqslant 2$. The latter is achieved by ensuring that every such boundary contains a topological copy of the unit interval or else a copy of the Cantor set whose subspace topology is homeomorphic to Kuratowski's graph topology, as defined in [4].

---

* Corresponding author.
  E-mail addresses: adow@uncc.edu (A. Dow), K.P.Hart@TUDelft.NL (K.P. Hart).
  URLs: http://www.math.uncc.edu/~adow, http://aw.twi.tudelft.nl/~hart.

The principal difference between our approach and that of [2] lies in the details of the constructions. In [2] the topology is introduced by way of resolutions; however, some of the arguments given in the paper need emending because, for example, Kuratowski's function does not have the properties asserted and used in Lemmas 2.2 and 2.3 of [2] respectively. We avoid this and use the Tietze–Urysohn theorem to extend Kuratowski's function to the whole plane and thus obtain, per Cantor set, a separable metric topology on the plane that extends the graph topology.[2]

Also, in [2] the construction of the Cantor sets is entwined with that of the topologies, which leads to some rather inaccessible lemmas. We separate the two strands and this, combined with the use of partial orders, leads to a cleaner and more perspicuous construction of the Cantor sets.

We begin, in Section 1, with a description of Kuratowski's function. We then show how to transplant the graph topology to an arbitrary Cantor set in the plane. The remainder of the paper is devoted to a recursive construction of the necessary Cantor sets and finishes with a verification of the properties of the new topology. An outline of the full construction can be found in Section 3.

1. Kuratowski's function

In this section we give a detailed description of Kuratowski's function ([4], see also [3, Exercise 1.2.E]) and the resulting topology on the Cantor set. We do this to make our note self-contained and because the construction makes explicit use of this description. We leave the verification of most of the properties to the reader.

Let C be the Cantor set, represented as the topological product $2^{\mathbb{N}}$, and for $x \in C$ write $\operatorname{supp} x = \{i\colon x(i) = 1\}$. We let D be the set of x for which $\operatorname{supp} x$ is finite, partitioned into the sets $D_k = \{x\colon |\operatorname{supp} x| = k\}$; put $k_x = |\operatorname{supp} x|$ and $N_x = \max \operatorname{supp} x$ for $x \in D$. Note that $D_0 = \{\mathbf{0}\}$, where $\mathbf{0}$ is the point with all coordinates 0. Let $E = C \setminus D$, the set of x for which $\operatorname{supp} x$ is infinite.

For $x \in C$ let c_x be the counting function of $\operatorname{supp} x$, so $\operatorname{dom} c_x = \{1, \ldots, k_x\}$ if $x \in D$ and $\operatorname{dom} c_x = \mathbb{N}$ if $x \in E$. Note that $N_{\mathbf{0}} = \operatorname{dom} c_{\mathbf{0}} = \emptyset$.

Now define
$$f(x) = \sum_{j \in \operatorname{dom} c_x} (-1)^{c_x(j)} 2^{-j}$$

Thus we use the parity of $c_x(j)$ to decide whether to add or subtract 2^{-j}. By convention an empty sum has the value 0, so $f(\mathbf{0}) = 0$.

Notation: if $x \in C$ and $n \in \mathbb{N}$ then $x \restriction n$ denotes the restriction of x to the set $\{1, 2, \ldots, n\}$. Also, $[x \restriction n]$ denotes the nth basic open set around x: $[x \restriction n] = \{y\colon y \restriction n = x \restriction n\}$.

For $x \in D$ we write $V_x = [x \restriction N_x]$. Using the V_x it is readily seen that the sets D_k are relatively discrete: simply observe that $V_x \cap \bigcup_{i \leqslant k_x} D_i = \{x\}$. In fact, for a fixed k the family $\mathcal{D}_k = \{V_x\colon x \in D_k\}$ is pairwise disjoint. For later use we put $D_x = \{y \in D_{k_x+1}\colon y \restriction N_x = x \restriction N_x\}$ and we observe that $V_x = \{x\} \cup \bigcup\{V_y\colon y \in D_x\}$.

1.1. Continuity

We begin by identifying the points of continuity of f.

Proposition 1. *The function f is continuous at every point of E.*

The function f is definitely not continuous at the points of D. This will become clear from the following discussion on the distribution of the values of f.

Proposition 2. *Let $t \in [-1, 1]$. The preimage $f^{\leftarrow}(t)$ is uncountable, crowded and its intersection with E is closed in E.*

Proposition 3. *Let $x \in D$ and $k = k_x$. Then x is an accumulation point of $f^{\leftarrow}(t)$ if and only if $f(x) - 2^{-k} \leqslant t \leqslant f(x) + 2^{-k}$.*

---

[2] M. Charalambous found a similar adaptation of Delistathis and Watson's construction with the added advantage that it works in ZFC alone.

1.2. The dimension of the graph

We identify f with its graph in $C \times [-1, 1]$ and we write $\mathbb{I} = [-1, 1]$. For $x \in D$ we let $I_x = [f(x) - 2^{-k_x}, f(x) + 2^{-k_x}]$. The discussion in the previous subsection can be summarized by saying that the closure of f in $C \times \mathbb{I}$ is equal to the set $K = f \cup \bigcup_{x \in D}(\{x\} \times I_x)$.

Proposition 4. $\operatorname{ind} f \leqslant 1$.

Proposition 5. *If $x \in E$ then* $\operatorname{ind}_{\langle x, f(x) \rangle} f = 0$.

Proposition 6. *If $x \in D$ then* $\operatorname{ind}_{\langle x, f(x) \rangle} f = 1$.

We put $\tau_f = \{O_f \colon O \text{ open in } C \times \mathbb{I}\}$, where $O_f = \{x \colon \langle x, f(x) \rangle \in O\}$; this is the topology of the graph, transplanted to C.

2. Making one Cantor set

We intend to copy the topology τ_f to many Cantor sets in the plane, or rather, we intend to construct many Cantor sets and copy τ_f to each of them. Here we describe how we will go about constructing just one Cantor set K, together with a homeomorphism $h \colon C \to K$, and how to refine the topology of the plane so that all points but those of $h[D]$ retain their usual neighbourhoods and so that at the points of $h[D]$ the dimension of K will be 1.

All we need to make a Cantor set are two maps $\sigma \colon D \to \mathbb{R}^2$ and $\ell \colon D \to \omega$. Using these we define $W(d) = B(\sigma(d), 2^{-\ell(d)})$ and $U(d) = B(\sigma(d), 2^{-\ell(d)-1})$ for each $d \in D$. We want the following conditions fulfilled:

(1) the sequence $\langle \sigma(e) \colon e \in D_d \rangle$ converges to $\sigma(d)$, for all d;
(2) $\operatorname{cl} W(e) \subseteq U(d) \setminus \{\sigma(d)\}$ whenever $e \in D_d$;
(3) $\{\operatorname{cl} W(d) \colon d \in D_n\}$ is pairwise disjoint for all n.

The following formula then defines a Cantor set:

$$K = \bigcap_{n=0}^{\infty} \operatorname{cl}\left(\bigcup \{W(d) \colon d \in D_n\}\right). \tag{‡}$$

One readily checks that $\{\sigma(d) \colon d \in D\}$ is a dense subset and that the map σ extends to a homeomorphism $h \colon C \to K$ with the property that $h[V_d] = K \cap W(d)$ for all $d \in D$. Also note that in (‡) we could have used the $U(d)$ instead of the $W(d)$ and that even $h[V_d] = K \cap U(d)$ for all d.

Copying the Kuratowski function from C to K is an easy matter: we let $f_K = f \circ h^{-1}$. To copy the topology τ_f to K *and* to preserve as much as possible of the Euclidean topology we use the sets $U(d)$ and $W(d)$.

We apply the Tietze–Urysohn theorem to extend f_K to a function $\bar{f}_K$ defined on the whole plane that is continuous everywhere except at the points of $\sigma[D]$. The topology τ_K that we get by identifying the plane with the graph of $\bar{f}_K$ is separable and metrizable and its restriction to K is the graph topology.

As will become clear below we cannot take just any extension of f_K because we will have to have some amount of continuity at the points of $\sigma[D]$. To this end we define for each $d \in D$ a closed set $F(d)$ by $F(d) = \operatorname{cl} U(d) \setminus \bigcup_{e \in D_d} W(e)$. The family $\{F(d) \colon d \in D\}$ is pairwise disjoint: if $F(d_1)$ and $F(d_2)$ meet then so do $U(d_1)$ and $U(d_2)$. Because of conditions (2) and (3) above this is only possible if, say, $U(d_1) \supseteq U(d_2)$. But, unless $d_1 = d_2$, this would entail $U(d_2) \subseteq U(e)$ for some $e \in D_d$ and so $F(d_2)$ would be disjoint from $F(d_1)$ after all.

The set $K^+ = K \cup \bigcup_{d \in D} F(d)$ is closed and we can extend f_K to K^+ by setting $f_K^+(x) = f(d)$, whenever $x \in F(d)$. Because for every $\varepsilon > 0$ there are only finitely many d for which the diameter of $F(d)$ is larger than ε this extended function is continuous at all points of $K \setminus \sigma[D]$. The new function f_K^+ is certainly continuous at the points of $K^+ \setminus K$ (it is even locally constant there), so we can apply the Tietze–Urysohn theorem to find a function $\bar{f}_K \colon \mathbb{R}^2 \to [-1, 1]$ that extends f_K^+ and that is continuous at all points, except those of $\sigma[D]$.

In fact, it not hard to verify that, if L is a subset of the plane that meets only finitely many of the sets $W(d)$ then the restriction of f_K to L is continuous. Indeed, we only have to worry about points in $\sigma[D]$. But if $d \in D$ then $F(d) \cap L$ contains a neighbourhood of $\sigma(d)$ in L and f_K^+ is constant on $F(d)$.

3. The plan

In this section we outline how we will construct a cosmic topology τ on a subset X of the plane that satisfies $\dim(X, \tau) = 1$ and $\operatorname{ind}(X, \tau) \geqslant 2$.

We let $\mathcal{Q}$ denote the family of all non-trivial line segments in the plane with rational end points. Our subset X will be $\mathbb{R}^2 \setminus A$, where $A = \{\langle p + \sqrt{2}, q \rangle: p, q \in \mathbb{Q}\}$. Note that A is countable, dense and disjoint from $\bigcup \mathcal{Q}$. Also note that, with respect to the Euclidean topology τ_e, one has $\operatorname{ind}(X, \tau_e) = 1$: on the one hand basic rectangles with end points in A have zero-dimensional boundaries (in X), so that $\operatorname{ind}(X, \tau_e) \leqslant 1$, and on the other hand, because X is connected we have $\operatorname{ind}(X, \tau_e) \geqslant 1$.

We will construct τ in such a way that its restrictions to $X \setminus \bigcup \mathcal{Q}$ and each element of $\mathcal{Q}$ will be the same as the restrictions of τ_e; this ensures that (X, τ) has a countable network: take a countable base $\mathcal{B}$ for the Euclidean topology of $X \setminus \bigcup \mathcal{Q}$, then $\mathcal{Q} \cup \mathcal{B}$ is a network for (X, τ). Also, the τ_e-interior of every open set in (X, τ) will be nonempty so that $\bigcup \mathcal{Q}$ and $X \setminus \bigcup \mathcal{Q}$ will be dense with respect to τ.

It what follows cl will be the closure operator with respect to τ and cl_e will be the Euclidean closure operator.

3.1. The topology

We let $\{(U_\alpha, V_\alpha): \alpha < \mathfrak{c}\}$ numerate all pairs of disjoint open sets in the plane whose union is dense and for each α we put $S_\alpha = \operatorname{cl}_e U_\alpha \cap \operatorname{cl}_e V_\alpha$. We shall construct for each α a Cantor set K_α in $X \cap S_\alpha$, unless there is a $Q_\alpha \in \mathcal{Q}$ that is contained in S_α. The construction of the K_α will be as described in Section 2, so that we will be able to extend τ_e to a topology τ_α whose restriction to K_α is a copy of the topology τ_f. For notational convenience we let I be the set of αs for which we have to construct K_α and for $\alpha \in \mathfrak{c} \setminus I$ we set $\tau_\alpha = \tau_e$. As an aside we mention that $\mathfrak{c} \setminus I$ is definitely not empty: if the boundary of U_α is a polygon with rational vertices then $\alpha \notin I$.

Thus we may (and will) define, for any subset J of $\mathfrak{c}$ a topology τ_J: the topology generated by the subbase $\bigcup_{\alpha \in J} \tau_\alpha$. The new topology τ will $\tau_\mathfrak{c}$.

There will be certain requirements to be met (the first was mentioned already):

(1) The restriction of τ to $X \setminus \bigcup \mathcal{Q}$ and each $Q \in \mathcal{Q}$ must be the same as that of the Euclidean topology;
(2) Different topologies must not interfere: the restriction of τ to K_α should be the same as that of τ_α;
(3) For each α, depending on the case that we are in, the set K_α or Q_α must be part of the τ-boundary of U_α.

If these requirements are met then the topology τ will be as required. We have already indicated that (1) implies that it has a countable network.

3.2. The inductive dimensions

To see that $\operatorname{ind}(X, \tau) \geqslant 2$ we take an element O of τ and show that its boundary is at least one-dimensional. There will be an α such that $\operatorname{cl}_e O = \operatorname{cl}_e U_\alpha$: there is $O' \in \tau_e$ such that $O \cap \bigcup \mathcal{Q} = O' \cap \bigcup \mathcal{Q}$ and we can take α such that $U_\alpha = \operatorname{int} \operatorname{cl}_e O'$ and $V_\alpha = \mathbb{R}^2 \setminus \operatorname{cl}_e U$. In case $\alpha \in I$ the combination of (2) and (3) shows that $\operatorname{ind} \operatorname{Fr} O \geqslant \operatorname{ind} K_\alpha = 1$ and in case $\alpha \notin I$ we use (1) and (3) to deduce that $\operatorname{ind} \operatorname{Fr} O \geqslant \operatorname{ind} Q_\alpha = 1$.

3.3. The covering dimension

As $\operatorname{ind}(X, \tau) \geqslant 2$ it is immediate that $\dim(X, \tau) \geqslant 1$. To see that $\dim(X, \tau) \leqslant 1$ we consider a finite open cover $\mathcal{O}$. Because (X, τ) is hereditarily Lindelöf we find that each element of $\mathcal{O}$ is the union of countably many basic open sets. This in turn implies that there is a countable set J such that $\mathcal{O} \subseteq \tau_J$. The topology τ_J is separable and metrizable and it will suffice to show that $\dim(X, \tau_J) \leqslant 1$.

If J is finite then we may apply the countable closed sum theorem: $O = X \setminus \bigcup_{\alpha \in J} K_\alpha$ is open, hence an F_σ-set, say $O = \bigcup_{i=1}^\infty F_i$. Each F_i is (at most) one-dimensional as is each K_α and hence so is X, as the union of countably many one-dimensional closed subspaces.

If J is infinite we numerate it as $\{\alpha_n : n \in \mathbb{N}\}$ and set $J_n = \{\alpha_i : i \leqslant n\}$. Then (X, τ_J) is the inverse limit of the sequence $\langle (X, \tau_{J_n}) : n \in \mathbb{N} \rangle$, where each bonding map $i_n : (X, \tau_{J_{n+1}}) \to (X, \tau_{J_n})$ is the identity. By Nagami's theorem ([6], see also [3, Theorem 1.13.4]) it follows that $\dim(X, \tau_J) \leqslant 1$.

4. The execution

The construction will be by recursion on $\alpha < \mathfrak{c}$. At stage α, if no Q_α can be found, we take our cue from Section 2 and construct maps $\sigma_\alpha : D \to S_\alpha$ and $\ell_\alpha : D \to \omega$, in order to use the associated balls $W_\alpha(d) = B(\sigma_\alpha(d), 2^{-\ell_\alpha(d)})$ in formula (‡) to make the Cantor set K_α. We also get a homeomorphism $h_\alpha : C \to K_\alpha$ as an extension of d_α and use this to copy Kuratowski's function to K_α: we set $f_\alpha = f \circ h_\alpha^{-1}$.

We use the procedure from the end of Section 2 to construct the topology τ_α. We let $U_\alpha(d) = B(\sigma_\alpha(d), 2^{-\ell_\alpha(d)-1})$, put $F_\alpha(d) = \operatorname{cl} U_\alpha(d) \setminus \bigcup_{e \in D_d} W_\alpha(e)$ and define K_α^+ and f_α^+ as above. We obtain τ_α as the graph topology from an extension $\bar{f}_\alpha$ of f_α^+.

4.1. The partial order

We construct σ_α and ℓ_α by an application of Martin's Axiom to a partial order that we describe in this subsection. To save on notation we suppress α for the time being. Thus, $S = S_\alpha$, $\sigma = \sigma_\alpha$, etc.

To begin we observe that $\bigcup \mathcal{Q} \cap S$ is dense in S: if $x \in S$ and $\varepsilon > 0$ then there are points a and b with rational coordinates in $B(x, \varepsilon)$ that belong to U and V respectively. The segment $Q = [a, b]$ belongs to $\mathcal{Q}$, is contained in $B(X, \varepsilon)$ and meets S. Actually, $Q \cap S$ is nowhere dense in Q because no subinterval of Q is contained in S—this is where we use the assumption that no element of $\mathcal{Q}$ is contained in S. There is therefore even a point y in $Q \cap S$ that belongs to $\operatorname{cl}(Q \cap U) \cap \operatorname{cl}(Q \cap V)$: orient Q so that a is its minimum, then $y = \inf(Q \cap V)$ is as required. It follows that the set S' of those $y \in S$ for which there is $Q \in \mathcal{Q}$ such that $y \in \operatorname{cl}(Q \cap U) \cap \operatorname{cl}(Q \cap V)$ is dense in S. We fix a countable dense subset T of S'. We also fix a numeration $\{a_n : n \in \mathbb{N}\}$ of A, the complement of our set X.

The elements p of our partial order $\mathbb{P}$ have four components:

(1) a finite partial function σ_p from D to T,
(2) a finite partial function ℓ_p from D to ω,
(3) a finite subset F_p of $\alpha \cap I$,
(4) a finite subset $\mathcal{Q}_p$ of $\mathcal{Q}$.

We require that $\operatorname{dom} \sigma_p = \operatorname{dom} \ell_p$ and we abbreviate this common domain as $\operatorname{dom} p$. It will be convenient to have $\operatorname{dom} p$ downward closed in D, by which we mean that if $e \in \operatorname{dom} p \cap D_d$ then $d \in \operatorname{dom} p$.

The intended interpretation of such a condition is that σ_p and ℓ_p approximate the maps σ and ℓ respectively; therefore we also write $W_p(d) = B(\sigma_p(d), 2^{-\ell_p(d)})$ and $U_p(d) = B(\sigma_p(d), 2^{-\ell_p(d)-1})$. The list of requirements in Section 2 must be translated into conditions that we can impose on σ_p and ℓ_p.

(1) $\|\sigma_p(e) - \sigma_p(d)\| < 2^{-N_e}$ whenever $d, e \in \operatorname{dom} p$ are such that $e \in D_d$, this will ensure that $\langle \sigma(e) : e \in D_d \rangle$ will converge to $\sigma(d)$;
(2) $\operatorname{cl}_e W_p(e) \subseteq U_p(d) \setminus \{\sigma_p(d)\}$ whenever $d, e \in \operatorname{dom} p$ are such that $e \in D_d$; and
(3) for every n the family $\{\operatorname{cl}_e W_p(d) : d \in D_n \cap \operatorname{dom} p\}$ is pairwise disjoint.

The order on $\mathbb{P}$ will be defined to make p force that for $\beta \in F_p$ and $Q \in \mathcal{Q}_p$ the intersection $\{\sigma(d) : d \in D\} \cap (K_\beta \cup Q)$ is contained in the range of σ_p, and even that when $d \notin \operatorname{dom} p$ the intersection $\operatorname{cl}_e W(d) \cap (K_\beta \cup Q)$ is empty. We also want p to guarantee that $K \cap \{a_i : i \leqslant |\operatorname{dom} p|\} = \emptyset$.

Before we define the order, however, we must introduce an assumption on our recursion that makes our density arguments go through with relatively little effort; unfortunately it involves a bit of notation.

For $x \in \bigcup \mathcal{Q}$ set $I_x = \{\beta \in I: x \in \sigma_\beta[D]\}$. For each $\beta \in I_x$ let $d_\beta = \sigma_\beta^\leftarrow(x)$ and write $D_{x,\beta} = D_{d_\beta}$. If it so happens that $q \in \mathbb{P}$ and $x = \sigma_q(d)$ for some $d \in D$ and if $e \in D_d \setminus \mathrm{dom}\, q$ then we must be able to choose an extension p of q with $e \in \mathrm{dom}\, p$, without interfering too much with the sets $W_\beta(a)$, where $\beta \in I_x$ and $a \in D_{x,\beta}$. The following assumption enables us to do this (and we will be able to propagate it):

(∗) If $x \in \bigcup \mathcal{Q}$ then for every finite subset F of $I_x \cap \alpha$ there is an $\varepsilon > 0$ such that the family $\mathcal{W}_{F,\varepsilon} = \{\mathrm{cl}_e\, W_\beta(a): \beta \in F, a \in D_{x,\beta}$ and $\sigma(a) \in B(x, \varepsilon)\}$ is pairwise disjoint.

It is an elementary exercise to verify that in such a case the difference $B(x, \varepsilon) \setminus \bigcup \mathcal{W}_{F,\varepsilon}$ is connected. Assumption (∗) will also be useful when we verify some of the properties of the topology τ.

We define $p \preccurlyeq q$ if

(1) σ_p extends σ_q and ℓ_p extends ℓ_q;
(2) $F_p \supseteq F_q$ and $\mathcal{Q}_p \supseteq \mathcal{Q}_q$;
(3) if $d \in \mathrm{dom}\, p \setminus \mathrm{dom}\, q$ and $i \leqslant |\mathrm{dom}\, q|$ then $a_i \notin \mathrm{cl}_e\, W_p(d)$;
(4) if $d \in \mathrm{dom}\, p \setminus \mathrm{dom}\, q$ and $J \in \mathcal{Q}_q \cup \{K_\beta: \beta \in F_p\}$ then $\mathrm{cl}_e\, W_p(d)$ is disjoint from J;
(5) if $d \in \mathrm{dom}\, q$ and $x = \sigma_q(d)$ and if $e \in \mathrm{dom}\, p \setminus \mathrm{dom}\, q$ is such that $e \in D_d$ then $\mathrm{cl}_e\, W_p(e)$ is disjoint from $\mathrm{cl}_e\, W_\beta(a)$ whenever $\beta \in F_q \cap A_x$ and $a \in D_{x,\beta}$.

It is clear that p and q are compatible whenever $\sigma_p = \sigma_q$ and $\ell_p = \ell_q$; as there are only countably many possible σs and ℓs we find that $\mathbb{P}$ is a σ-centered partial order.

4.2. Dense sets

In order to apply Martin's Axiom we need, of course, a suitable family of dense sets.

For $\beta < \alpha$ the set $\{p: \beta \in F_p\}$ is dense.   Given p and β extend p by adding β to F_p.

For $Q \in \mathcal{Q}$ the set $\{p: Q \in \mathcal{Q}_p\}$ is dense.   Given p and Q extend p by adding Q to $\mathcal{Q}_p$.

For $n \in \mathbb{N}$ the set $\{p: |\mathrm{dom}\, p| \geqslant n\}$ is dense.   This follows from the density of the sets below.

For $e \in D$ the set $\{p: e \in \mathrm{dom}\, p\}$ is dense.   Here is where we use assumption (∗). Since every $e \in D$ has only finitely many predecessors with respect to the transitive closure of the relation "$D_d \ni e$" it will suffice to consider the case where $q \in \mathbb{P}$ and $e \in D_d \setminus \mathrm{dom}\, q$ for some $d \in \mathrm{dom}\, q$.

We extend q to a condition p by setting $F_p = F_q$, $\mathcal{Q}_p = \mathcal{Q}_q$, $\mathrm{dom}\, p = \{e\} \cup \mathrm{dom}\, q$ and by defining $d_p(e)$ and $\ell_p(e)$ as follows. Let $x = \sigma_q(d)$, put $n = k_e$ and consider $H = \bigcup\{\mathrm{cl}_e\, W_q(a): a \in D_{n+1} \cap \mathrm{dom}\, p \cap D_d\}$.

Fix $\varepsilon_1 \leqslant 2^{-N_e}$ so that $B(x, 2\varepsilon_1)$ is disjoint from H, this is possible because of condition (2) in the definition of the elements of $\mathbb{P}$. Observe that if we choose $\sigma_p(e)$ and $\ell_p(e)$ in such a way that $\mathrm{cl}_e\, W_p(e) \subseteq B(x, \varepsilon_1/2)$ then p is an element of $\mathbb{P}$.

Next, using (∗), find $\varepsilon_2 \leqslant \varepsilon_1/2$ that works for the finite set $F_q \cap I_x$. The set $W = \{x\} \cup \bigcup \mathcal{W}_{F,\varepsilon_2}$ is closed and does not separate the ball $B(x, \varepsilon_2)$ and the set S *does* separate this ball because the latter meets both U and V. Therefore we can find a point y in $S \cap B(x, \varepsilon_2) \setminus W$; we choose $\delta > 0$ so small that $\mathrm{cl}_e\, B(y, \delta) \subseteq B(x, \varepsilon_2) \setminus W$.

The set $S \cap B(y, \delta)$ separates $B(y, \delta)$, hence it is (at least) one-dimensional. The union of the K_β (for $\beta \in F_q$) together with the $Q \cap S$ (for $Q \in \mathcal{Q}_q$) is zero-dimensional because each individual set is: each K_β is a Cantor set and each $Q \cap S$ is nowhere dense in Q and hence zero-dimensional. This means that, finally, we can choose $\sigma_p(e)$ in $T \cap B(y, \delta)$ but not in this union and then we take $\ell_p(e)$ so large that $\mathrm{cl}_e\, W_p(e)$ is a subset of $B(y, \delta)$ minus that union. Also, at this point we ensure that $a_i \notin \mathrm{cl}_e\, W_p(e)$ for $i \leqslant |\mathrm{dom}\, q|$: this is possible because $\sigma_p(e) \notin A$.

We have chosen $W_p(e)$ to meet requirements (3), (4) and (5) in the definition of $p \preccurlyeq q$.

4.3. A generic filter

Let G be a filter on $\mathbb{P}$ that meets all of the above dense sets. Then $\sigma_\alpha = \bigcup\{\sigma_p \colon p \in G\}$ and $\ell_\alpha = \bigcup\{\ell_p \colon p \in G\}$ are the sought after maps. We define W_α and K_α as in Section 2.

Assumption $(*)$ *is propagated.* In verifying this we only have to worry about the points in $\sigma_\alpha[D]$ of course.

Therefore let $x \in \sigma_\alpha[D]$ and let F be a finite subset of $I_x \cap \alpha$; we have to find an ε for $F' = F \cup \{\alpha\}$. First fix ε_1 that works for F itself. Next take $p \in G$ such that $d_\alpha \in \operatorname{dom} p$ and $F \subseteq F_p$. Using condition (5) in the definition of $\preccurlyeq$ and a density argument we find that $\operatorname{cl}_e W_\alpha(e)$ is disjoint from $\operatorname{cl}_e W_\beta(a)$ whenever $e \in D_{x,\alpha} \setminus \operatorname{dom} p$, $\beta \in F$ and $a \in D_{x,\beta}$. Now choose ε smaller than ε_1 and all distances $\|x - \sigma_\alpha(e)\|$, where $e \in D_{x,\alpha} \cap \operatorname{dom} p$. Then $\mathcal{W}_{F',\varepsilon}$ is pairwise disjoint.

K_α *meets* K_β *in a finite set whenever* $\beta < \alpha$. Let $\beta \in \alpha \cap I$ and take $p \in G$ such that $\beta \in F_p$. Choose n such that $\operatorname{dom} p \subseteq \bigcup_{k \leqslant n} D_k$. By formula (‡) we know that $K_\alpha \subseteq \operatorname{cl}_e(\bigcup\{W_d \colon d \in D_{n+1}\})$ the latter closure is equal to $\bigcup_{k \leqslant n} D_k \cup \bigcup\{W_d \colon d \in D_{n+1}\}$ and the intersection of *this* set with K_β is contained in $\operatorname{dom} p$; this follows from condition (4) in the definition of $\preccurlyeq$.

K_α *meets each* $Q \in \mathcal{Q}$ *in a finite set.* The proof is identical to the previous one: take $p \in G$ with $Q \in \mathcal{Q}$.

5. The remaining properties of the topologies

We check conditions (1), (2) and (3) from Section 3.

A useful observation is that a typical new basic neighbourhood of a point x of $\bigcup \mathcal{Q}$ contains a set of the form $O(x, \varepsilon, G) = B(x, \varepsilon) \cap \bigcap_{\beta \in G} F_\beta(d_\beta)$, where $\varepsilon > 0$ and G is a finite subset of I_x.

5.1. $X \setminus \bigcup \mathcal{Q}$ retains its Euclidean topology

This is immediate from the observation that every function $\bar{f}_\alpha$ (for $\alpha \in I$) is continuous at the points of $X \setminus \bigcup \mathcal{Q}$.

5.2. Each $Q \in \mathcal{Q}$ retains its Euclidean topology

We should show that $\bar{f}_\alpha \restriction Q$ is continuous for each α in I and each $Q \in \mathcal{Q}$. The only points at which this restriction could possibly be discontinuous are those in $\sigma_\alpha[D] \cap Q$, which is a finite set. Let $d \in D$ be such that $x = \sigma_\alpha(d) \in Q$. By construction all but finitely many of the sets $\operatorname{cl} W_\alpha(e)$, where $e \in D_d$, meet Q. This implies that $F_\alpha(d) \cap Q$ is actually a neighbourhood of x in Q. As $\bar{f}_\alpha$ is constant on $F_\alpha(d)$ this shows that $\bar{f}_\alpha \restriction Q$ is continuous at x.

5.3. τ_α and τ_β do not interfere

If $\alpha \neq \beta$ then there are only finitely many points in $K_\alpha \cap K_\beta$ and it is only at these points that τ_α and τ_β might interfere and even then only at a point of $\sigma_\alpha[D] \cap \sigma_\beta[D]$. Let x be such a point and apply assumption $(*)$ to the set $G = \{\alpha, \beta\}$ to find $\varepsilon > 0$ such that $\mathcal{W}_{G,\varepsilon}$ is pairwise disjoint. But then $f_\alpha^+ \restriction K_\beta$ is constant on a neighbourhood of x in K_β, namely $O(x, \varepsilon, \{\alpha\}) \cap K_\beta$ and, by symmetry, $f_\beta^+ \restriction K_\alpha$ is constant on the neighbourhood $O(x, \varepsilon, \{\beta\}) \cap K_\alpha$ of x in K_α.

5.4. Q_α is still in the boundary of U_α

If $x \in Q_\alpha$ then, by construction, all points of the intersection $Q_\alpha \cap O(x, \varepsilon, G)$ (except x itself) belong to the Euclidean interior of $O(x, \varepsilon, G)$, Because these points are in the boundary of U_α that interior meets both U_α and V_α. Therefore each basic neighbourhood of x meets these sets as well.

5.5. K_α is still in the boundary of U_α

Let $x \in K_\alpha$, assume $I_x \neq \emptyset$ and consider some $O(x, \varepsilon, G)$.

If $\alpha \notin I_x$ then the same argument as above will work: the intersections $B(x, \varepsilon) \cap K_\alpha$ and $O(x, \varepsilon, G) \cap K_\alpha$ are equal when ε is small enough.

If $\alpha \in I_x$ then we assume $\alpha \in G$ and observe that if ε is small enough then $O(x, \varepsilon, G)$ is a Euclidean neighbourhood of many points of S_α.

References

[1] A.V. Arhangelskij, A survey of some recent advances in general topology, old and new problems, in: Actes du Congrès International des Mathématiciens (Nice, 1970), Tome 2, 1971, pp. 19–26, MR 0428244 (55 #1269).

[2] G. Delistathis, S. Watson, A regular space with a countable network and different dimensions, Transactions of the American Mathematical Society 352 (9) (2000) 4095–4111, MR 1661301 (2000m:54039).

[3] R. Engelking, Theory of Dimensions Finite and Infinite, Sigma Series in Pure Mathematics, vol. 10, Heldermann Verlag, Lemgo, 1995, MR 1363947 (97j:54033).

[4] C. Kuratowski, Une application des images de fonctions à la construction de certains ensembles singuliers, Mathematica 6 (1932) 120–123, Zbl 0005.05701, JFM 58.0627.04.

[5] E. Michael, $\aleph_0$-spaces, Journal of Mathematics and Mechanics 15 (1966) 983–1002, MR 0206907 (34 #6723).

[6] K. Nagami, Finite-to-one closed mappings and dimension. II, Proceedings of the Japan Academy 35 (1959) 437–439, MR 0113214 (22 #4052).

[7] N. Vedenissoff, Remarques sur la dimension des espaces topologiques, Uchenye Zapiski Moskov. Gos. Univ. Matematika 30 (1939) 131–140 (in Russian, with French summary), MR0002526 (2,69i).