MATH 452

General Topology

2007

§0. Preliminary remarks.

General Topology, sometimes called Analytic Topology, is the most basic part of topology. There are other subjects like Geometric Topology, Algebraic Topology, and Differential Topology, which exploit some of the ideas we shall meet; and much of Functional Analysis needs others. Broadly speaking, these further subjects deal with more specialized spaces. My aim will be, therefore, to present an array of facts, ideas, and definitions that are useful in many contexts, rather than to go very deeply into any of them. Since the theory arose in several different contexts before it was unified as an independent topic, there is an amazing amount of terminology that is in constant use despite being to some extent superfluous; I shall attempt to accustom you to the words that are most commonly encountered, even though many of them could be circumvented in a strictly logical exposition. With one exception, I shall not prove any of the really difficult theorems. Sometimes I shall just assert results with a proof sign; this means "prove it yourself", even if I do not set it as an assignment exercise.

There is also a slight difficulty: because *general* topology is so widely relevant, it is almost inevitable that there will be some duplication of material from other courses. For instance, I shall have to say something about the axiom of choice (which, be warned, I shall casually assume throughout, but at some points I shall need in a non-trivial way) and about metric spaces. I may also mention something about categories. And I make no apology for the fact that the exposition is rather dry; until quite late in the course, it is a matter of getting used to the concepts rather than of proving anything really striking.

The subject was developed with some enthusiasm from about 1922 to about 1955, and is still active in some quarters and in some flavours. By 1970 or so it was generally regarded as a little démodé, and I am told Dieudonné, who was given to dismissive remarks, commented in private that it is only interesting as a necessary background to the theory of topological vector spaces. Interesting or not, there are several non-trivial books about it. A relatively accessible one is by Kelley; it has both advantages and disadvantages (he devotes a lot of space to his own work), but is rather clearly written. A more modern and far more complete one is by Engelking, which, despite or because of being comparatively encyclopædic, is not a book for coherent reading. Indeed it is heavy going however you try to approach it, but it contains a great deal of information. There are several more elementary books, one of which (by Munkres) has been used as a textbook for this course.

§1. Topologies.

Definition 1.1. Let Ω be any set. A *topology* on, or in, Ω is a class $\mathcal{G}$ of subsets of Ω satisfying the three following properties.

(a) $\emptyset \in \mathcal{G}$ and $\Omega \in \mathcal{G}$.

(b) If $U \in \mathcal{G}$ and $V \in \mathcal{G}$, then $U \cap V \in \mathcal{G}$.

(c) If $\mathcal{U} \subseteq \mathcal{G}$, then $\bigcup_{U \in \mathcal{U}} V \in \mathcal{G}$.

Recall that the "power class" of Ω is the set $\mathcal{P}(\Omega)$ of all subsets of Ω; thus $\mathcal{P}(\{1,2\})$, for instance, is $\{\emptyset, \{1\}, \{2\}, \{1,2\}\}$). Hence, a topology on Ω is a subclass $\mathcal{G}$ of $\mathcal{P}(\Omega)$ with certain special properties.

It is customary to describe the *members* of $\mathcal{G}$, each of which is a *subset* of Ω, as *open sets* of Ω (with respect to the topology $\mathcal{G}$). Topologies are also often denoted by the letter $\mathcal{T}$ or something similar.

By a *topological space*, we mean an ordered pair $(\Omega, \mathcal{G})$ consisting of a set and a topology thereon, and we shall sometimes use this ordered pair notation to avoid ambiguity. Mostly, however, only one topology on a given set Ω interests us, and we may then speak of "open sets" without mentioning $\mathcal{G}$ explicitly, and Ω, with the agreed topology understood in the background, may itself be described as a "topological space". Then, the three axioms above may be stated informally as:

the empty set and the whole space are open; the intersection of two open sets is also open; and the union of any class of open sets is open.

As with vector spaces (and others), one may sometimes speak simply of a "space" if it is clear that it refers to a topological space.

Where *(c)* is concerned, notice that $\bigcup_{V \in \emptyset} V = \emptyset$. (It is the set consisting of all the elements of all the members of the class $\emptyset$, i.e. of all the members of nothing at all.) And, crudely speaking, it does not matter how many open sets there are in the class $\mathcal{U}$; the union of *any number* (finite or "infinite", and no matter how "infinite") of open sets must still be an open set. The class $\mathcal{U}$ may be uncountable, or countable, or finite.

Lemma 1.2. *If $\mathcal{G}$ is a topology on Ω, $n \in \mathbb{N}$, and $U_1, U_2, \dots, U_n$ are open sets in Ω with respect to $\mathcal{G}$, then $U_1 \cap U_2 \cap \dots \cap U_n$ is also open with respect to $\mathcal{G}$.*

Proof. Use 1.1*(b)* and induction on n. $\square$

It is very important to grasp that mathematical induction only proves the result to hold for an arbitrary natural number n. For an infinite sequence $U_1, U_2, U_3, \dots$ of open sets, the intersection $\bigcap_{k=1}^{\infty} U_k$ need not be open, though the Lemma ensures that $\bigcap_{k=1}^{n} U_k$ is open for each finite n.

By contrast, $\bigcup_{k=1}^{\infty} U_k$ must be open by 1.1*(c)*. As I remarked above, the union of a *finite* class of open sets, as in $\bigcup_{k=1}^{n} U_k$, is open, and the union of a *countable* class of open sets, as in $\bigcup_{k=1}^{\infty} U_k$, is open. (Recall that a class is "countable", or more precisely "countably infinite", if its elements may be set out in an infinite sequence, such as $U_1, U_2, U_3, \dots$ in this case.) But axiom 1.1*(c)* applies to *any* subclass $\mathcal{U}$ of $\mathcal{G}$, no matter how large; and $\mathcal{U}$ may be massively uncountable. There is no problem in defining the union of *any* subclass of $\mathcal{P}(\Omega)$. A silly (because rather obvious) example is that the intervals $(1, x)$ in $\mathbb{R}$ are open sets in the usual topology, soon to be defined, on $\mathbb{R}$; there are uncountably many of them, but the axiom says $\bigcup_{x \in \mathbb{R}} (1, x)$ must also be open in $\mathbb{R}$.

Definition 1.3. *(a)* Let Ω be any set whatever. If $\mathcal{G}_1, \mathcal{G}_2$ are topologies on Ω, we say that $\mathcal{G}_1$ is *larger* or *finer* or *stronger* than $\mathcal{G}_2$ if $\mathcal{G}_1 \supseteq \mathcal{G}_2$ (as subclasses of $\mathcal{P}(\Omega)$); that is, if every subset of Ω that is open with respect to $\mathcal{G}_2$ is also open with respect to $\mathcal{G}_1$. In that case we also say that $\mathcal{G}_2$ is *smaller* or *coarser* or *weaker* than $\mathcal{G}_1$. People sometimes write $\mathcal{G}_1 \geq \mathcal{G}_2$. It is easy to see that this is a partial order on the set of topologies on Ω.

(b) The *discrete* topology on Ω is that topology with respect to which every subset of Ω is open; that is, $\mathcal{G} := \mathcal{P}(\Omega)$.

(c) The *indiscrete* topology on Ω is that topology with respect to which the only open sets are $\emptyset$ and Ω; that is, $\mathcal{G} := \{\emptyset, \Omega\}$.

The discrete topology is the finest possible topology on Ω; it is finer than any other topology. In view of 1.1*(c)*, it is characterized (i.e. is completely determined) by the fact that all singletons are open with respect to it. We often say “par abus de langage” that in the discrete topology *points are open*.

The indiscrete topology is, likewise, coarser than any other topology on Ω.

These two topologies are uninteresting. Any statement about them amounts to a statement about set theory and nothing more. But they are useful as examples.

Definition 1.4. *(a)* The *cofinite topology* on Ω is the topology whose members (i.e. the open sets) are $\emptyset$ itself and all those subsets of Ω whose complements are finite.

(b) The *cocountable* topology on Ω is the topology whose open sets are $\emptyset$ itself and all those subsets of Ω whose complements are finite or countable.

It is easy to see that these are indeed topologies on Ω, and that the cofinite topology is coarser than the cocountable topology. If Ω is finite, the cofinite topology is just the discrete topology; if Ω is countably infinite, the cocountable topology is the discrete topology and is strictly finer than the cofinite topology. If Ω is a singleton (or empty), both these topologies coincide with the indiscrete topology as well.

Example 1.5. Let Ω be any set with at least two elements, and suppose $a \in \Omega$. Let

$$\mathcal{G}_\Omega := \{\Omega, \emptyset, \mathcal{P}(\Omega \setminus \{a\})\}.$$

Then $\mathcal{G}_\Omega$ is a topology on Ω.

To specify a topology, we should in principle describe all the open sets. However, there may be less explicit ways of achieving this.

Definition 1.6. *(a)* Let $\mathcal{G}$ be a topology on Ω. A *base* for $\mathcal{G}$ is a subclass $\mathcal{B} \subseteq \mathcal{G}$ such that every nonempty member of $\mathcal{G}$ is a union of members of $\mathcal{B}$. (In other words, $\mathcal{B}$ is a class of open sets such that every open set is a union of *some* of the members of $\mathcal{B}$. The empty set $\emptyset$ is, of course, the union of the sets in the empty subclass of $\mathcal{B}$.)

(b) The topology $\mathcal{G}$ (or the topological space Ω) satisfies the *second axiom of countability* (or, vulgarly, is *second countable*) if it has a countable base.

The discrete topology has a very simple base, namely the class of all singletons. It is a base because, if $\emptyset \neq A \in \mathcal{P}(\Omega)$, then $A = \bigcup_{a \in A} \{a\}$, a union of singletons.

There are many other possible bases; indeed, it is clear that in this case any class of subsets that contains all the singletons is a base, and that the class of singletons is the smallest possible base. (Usually, there will be many bases for a given topology, but there will not be a smallest possible base.)

It is rather exceptional for a topology to be itself countable (for there to be only countably many open sets in all), but second countability means that it can be built out of countably many "building bricks".

The question arises whether a base for a topology must have any special properties (in the way that a topology itself is not just an arbitrary subclass of $\mathcal{P}(\Omega)$).

Lemma 1.7. *Let $\mathcal{G}$ be a topology on the set Ω, and let $\mathcal{B}$ be a base for $\mathcal{G}$. Then*

(a) *if $B_1, B_2 \in \mathcal{B}$ and $x \in B_1 \cap B_2$, there exists $B \in \mathcal{B}$ such that $x \in B \subseteq B_1 \cap B_2$;*

(b) *for any $x \in \Omega$, there is some $B_0 \in \mathcal{B}$ such that $x \in B_0$.*

Proof. *(a)* $B_1 \cap B_2$ is open by 1.1*(b)*, so it is the union of a subclass $\mathcal{A}$ of $\mathcal{B}$. Then $x \in B_1 \cap B_2 = \bigcup_{A \in \mathcal{A}} A$, and there exists $B \in \mathcal{A}$ such that $x \in B \subseteq B_1 \cap B_2$.

(b) Since Ω is itself open by 1.1*(a)*, it is a union of members of $\mathcal{B}$. ☐

Lemma 1.8. *Let $\mathcal{B}$ be a class of subsets of Ω enjoying the properties* 1.7(a), (b). *Then there is a unique topology on Ω for which $\mathcal{B}$ is a base.*

Proof. Let $\mathcal{G} := \{\bigcup_{A \in \mathcal{A}} A : \mathcal{A} \subseteq \mathcal{B}\}$, the class of subsets of Ω that can be expressed as unions of sets belonging to $\mathcal{B}$. We wish to show that $\mathcal{G}$ is a topology in Ω.

Firstly, $\emptyset = \bigcup_{A \in \emptyset} A \in \mathcal{G}$, as $\emptyset \subseteq \mathcal{B}$. And $\Omega = \bigcup_{A \in \mathcal{B}} A \in \mathcal{G}$ by 1.7*(b)*.

Secondly, let $U, V \in \mathcal{G}$. There are subclasses $\mathcal{U}, \mathcal{V}$ of $\mathcal{B}$ such that $U = \bigcup_{P \in \mathcal{U}} P$ and $V = \bigcup_{Q \in \mathcal{V}} Q$. Hence, if $x \in U \cap V$, there must be some $P \in \mathcal{U}$ and some $Q \in \mathcal{V}$ such that $x \in P$ and $x \in Q$; then $x \in P \cap Q$, and 1.7*(a)* tells us there is some $B(x) \in \mathcal{B}$ such that $x \in B(x) \subseteq P \cap Q \subseteq U \cap V$. For each x, we *choose* a suitable $B(x)$. Now $W := \bigcup_{x \in U \cap V} B(x)$ is a union of members of $\mathcal{B}$, so belongs to $\mathcal{G}$. $B(x) \subseteq U \cap V$ for each x, so $W \subseteq U \cap V$. For $x \in U \cap V$, $x \in B(x) \subseteq W$; thus $U \cap V \subseteq W$. Hence, $U \cap V = W \in \mathcal{G}$. [Notice that we have used the axiom of choice here.]

Finally, any union of sets from $\mathcal{G}$ is a union of unions of sets from $\mathcal{B}$, so is itself a union of sets from $\mathcal{B}$.

Hence, the three axioms for a topology (see 1.1) are satisfied by $\mathcal{G}$. It is evident that $\mathcal{B}$ is a base for $\mathcal{G}$, and that $\mathcal{G}$ is the only possible topology for which $\mathcal{B}$ is a base. ☐

We may say that $\mathcal{G}$ is the topology *generated* by the base $\mathcal{B}$.

Remark 1.9. *(a)* Suppose that $\mathcal{B}_\epsilon$ is a base for the topology $\mathcal{G}_\epsilon$, for $\epsilon = \pm 1$. If $\mathcal{B}_{-1} \subseteq \mathcal{B}_1$, then $\mathcal{G}_{-1} \subseteq \mathcal{G}_1$. Furthermore, $\mathcal{G}_1 = \mathcal{G}_{-1}$ if and only if, for $\epsilon = \pm 1$, each member of $\mathcal{B}_\epsilon$ is a union of members of $\mathcal{B}_{-\epsilon}$.

(b) Condition 1.7*(a)* is satisfied if, for any $B_1, B_2 \in \mathcal{B}$, either $B_1 \cap B_2 \in \mathcal{B}$ or $B_1 \cap B_2 = \emptyset$. ($\mathcal{B}$ may well not contain $\emptyset$, but will often contain two disjoint sets.) This condition is quite often satisfied, but should not be assumed in general.

(c) In the proof above, I showed from 1.7*(a)* (assuming the axiom of choice) that $B_1 \cap B_2$ is itself a union of members of $\mathcal{B}$; this assertion in turn obviously implies

1.7*(a)*. We could avoid the axiom of choice, by modifying 1.7*(a)* to the requirement that $B_1 \cap B_2$ be a union of members of $\mathcal{B}$.

Definition 1.10. The *standard topology* (or Euclidean topology, or "usual" topology) on $\mathbb{R}$ is the topology that has the class of open intervals $\{(a,b) : ab \in \mathbb{R}\}$ as a base. Here $(a,b) = \{x \in \mathbb{R} : a < x < b\}$. (Thus $(a,b) = \emptyset$ when $b \le a$.)

In this example 1.9*(b)* applies.

Definition 1.11. A *metric* on the set Ω is a function $d : \Omega \times \Omega \longrightarrow \mathbb{R}$ such that, for any $x, y, z \in \Omega$,

(a) $d(x,z) \le d(x,y) + d(z,y)$, [this is the *triangle inequality*]

(b) $d(x,y) = 0$ if and only if $x = y$.

If *(b)* is weakened to require only that $d(x,x) = 0$ for each $x \in \Omega$, d is a *pseudometric*. It is easy to deduce that a pseudometric must take nonnegative values and must be symmetric: $d(x,y) = d(y,x)$ for all $x, y \in \Omega$.

A *(pseudo)metric space* is a pair (Ω, d) consisting of a set Ω and a (pseudo)metric d on Ω. One may say that Ω is a *(pseudo)metric space with (pseudo)metric* d.

If $x \in \Omega$ and $r \in \mathbb{R}$, the ((pseudo)metric) *ball about x of radius r with respect to the (pseudo)metric* d is the set $B_d(x;r) = B(x;r) := \{y \in \Omega : d(x,y) < r\}$.

$B(x;r)$ is often called the *open* ball about x of radius r. It is empty unless $r > 0$, so its definition is often stated with the requirement that r should be positive.

Lemma 1.12. *The class of balls in a (pseudo)metric space (Ω, d) satisfies* 1.7(a), (b).

Proof. Suppose that $x \in B(x_1;r_1) \cap B(x_2;r_2)$. Then $d(x,x_i) < r_i$ for $i = 1, 2$. Take $r := \min(r_1 - d(x,x_1), r_2 - d(x,x_2))$, which is the lesser of two positive numbers and itself positive. I show that $B(x;r) \subseteq B(x_1;r_1) \cap B(x_2;r_2)$.

Suppose that $y \in B(x;r)$. Then, for each i,

$$d(x_i, y) \le d(x,y) + d(x,x_i) < r + d(x,x_i) = r_i\,;$$

thus $y \in B(x_i;r_i)$ for each i, and the assertion is proved. This demonstrates 1.7*(a)*.

If $x \in \Omega$, then $x \in B(x;r)$ for any $r > 0$. This establishes 1.7*(b)*. □

Definition 1.13. The topology on the metric space (Ω, d) generated by the class of metric balls is called the *metric topology* on Ω defined by the metric d.

In most undergraduate courses, the concept of a metric space is introduced before that of a topology, and a set U is defined to be open with respect to the metric if, for every $x \in U$, there is some $r > 0$ such that $B(x;r) \subseteq U$. (The number r will in principle depend on x.) This gives the same open sets as 1.12.

Metric spaces were invented by Fréchet in 1906. There were several attempts to define topological spaces in the years following; the first really satisfactory definition (not 1.1, but equivalent) was by Kuratowski in 1922. But it is probably true that most of the examples of topological spaces that people were thinking of at the time were metric spaces, and it remains true that those examples are of great interest.

In $\mathbb{R}^n$ and $\mathbb{C}^n$, there are the standard "Euclidean" or "Hermitian" metrics:

$$d((x_1, x_2, \ldots, x_n), (y_1, y_2, \ldots, y_n)) = \sqrt{\sum_{i=1}^{n} |x_i - y_i|^2} \,.$$

(Here x_i, y_i are real numbers for $\mathbb{R}^n$ and complex numbers for $\mathbb{C}^n$, and $|\,|$ denotes the absolute value or the modulus accordingly; that 1.11 is satisfied is a standard consequence of the Cauchy-Schwarz inequality.) In particular, when $n = 1$ the topology on $\mathbb{R}$ defined from the standard metric coincides with that described at 1.10.

We have seen that a subclass $\mathcal{B}$ of $\mathcal{P}(\Omega)$ can only be a base for a topology if it satisfies 1.7*(a)*, *(b)*. But suppose $\mathcal{S}$ is an *arbitrary* subclass of $\mathcal{P}(\Omega)$. Is it possible for $\mathcal{S}$ to be a subclass, not perhaps a base, of some topology? Another way of looking at this question is to ask whether there are any subclasses of $\mathcal{P}(\Omega)$ that are disqualified by their "internal structure" from consisting of open sets in some topology.

Remark 1.14. Suppose that $\mathcal{S}$ is some class of open sets in a topology $\mathcal{G}$, $\mathcal{S} \subseteq \mathcal{G}$. Then all finite intersections of sets in $\mathcal{S}$ must also be open:

$$\mathcal{S} \subseteq \mathcal{S}^{\cap} := \{S_1 \cap S_2 \cap \cdots \cap S_n : n \in \mathbb{N} \;\&\; S_1, S_2, \ldots, S_n \in \mathcal{S}\} \subseteq \mathcal{G}\,, \qquad (1)$$

and so must be all unions of finite intersections of sets in $\mathcal{S}^{\cap}$:

$$\mathcal{S}^{\cap} \subseteq \widehat{\mathcal{S}} := \left\{\bigcup_{A \in \mathcal{A}} A : \mathcal{A} \subseteq \mathcal{S}^{\cap}\right\} \subseteq \mathcal{G}\,. \qquad (2)$$

If, on the other hand, $\mathcal{S}$ is *any* class of subsets of Ω, $\mathcal{S} \subseteq \mathcal{P}(\Omega)$, $\widehat{\mathcal{S}}$ may be defined by (1) and (2), and clearly satisfies 1.1*(b)* and 1.1*(c)*; and $\emptyset \in \widehat{\mathcal{S}}$ because we may take $\mathcal{A}$ to be $\emptyset$ in (2). Clearly, then, $\widehat{\mathcal{S}} \cup \{\Omega\}$ is a topology in Ω, for *any* class $\mathcal{S}$ of subsets of Ω. Furthermore, it is the coarsest topology for which $\mathcal{S}$ consists of open sets. We may say it is the topology *generated* by $\mathcal{S}$.

Definition 1.15. A class $\mathcal{S}$ of subsets of Ω is a *subbase* for the topology $\mathcal{G}$ on Ω if $\widehat{\mathcal{S}} = \mathcal{G}$; that is, if every open set of $\mathcal{G}$ is a union of finite intersections of sets of $\mathcal{S}$.

Lemma 1.16. *The class $\mathcal{S}$ is a subbase for the topology it generates if and only if the union of the members of $\mathcal{S}$ is the whole of Ω. Then $\mathcal{S}^{\cap}$ is a base for the topology.* ☐

§2. Other ways of specifying topologies.

The modern definition 1.1 is only one of several equivalent formulations; indeed, the original definition of Kuratowski in 1922 was the one I am about to give, in terms of a "closure operation". These various other definitions of a topology are of no great importance in themselves, but their associated vocabulary is in common use, although, logically speaking, we shall just be dressing the same concepts in different words.

Definition 2.1. A *closure operation* in a set Ω is a mapping $\mathrm{cl} : \mathcal{P}(\Omega) \longrightarrow \mathcal{P}(\Omega)$ such that the following properties are satisfied.

(a) $\mathrm{cl}(\emptyset) = \emptyset$.

(b) For every $A \in \mathcal{P}(\Omega)$, $A \subseteq \mathrm{cl}(A)$.

(c) For any $A, B \in \mathcal{P}(\Omega)$, $\mathrm{cl}(A \cup B) = \mathrm{cl}(A) \cup \mathrm{cl}(B)$.

(d) For any $A \in \mathcal{P}(\Omega)$, $\mathrm{cl}(\mathrm{cl}(A)) = \mathrm{cl}(A)$.

Given such a closure operation cl in Ω, describe $A \in \mathcal{P}(\Omega)$ as a *closed set of* cl (or as *closed with respect to* cl) if $\mathrm{cl}(A) = A$.

Lemma 2.2. *Let* cl *be a closure operation in* Ω*, and let* $\mathcal{F}$ *be the class of subsets of* Ω *that are closed with respect to* cl *in the sense just defined. Then*

(i) $\emptyset \in \mathcal{F}$ *and* $\Omega \in \mathcal{F}$,

(ii) *if* $F_1, F_2 \in \mathcal{F}$, *then* $F_1 \cup F_2 \in \mathcal{F}$,

(iii) *if* $\mathcal{Q}$ *is any subclass of* $\mathcal{F}$*, then* $\bigcap_{Q \in \mathcal{Q}} Q$ *belongs to* $\mathcal{F}$*.*

Proof. 2.1*(a)* shows that $\emptyset$ is closed. By 2.1*(b),* $\Omega \subseteq \mathrm{cl}(\Omega) \in \mathcal{P}(\Omega)$, so $\Omega = \mathrm{cl}(\Omega)$; thus Ω is closed. *(ii)* follows instantly from 2.1*(c)*.

Suppose in 2.1*(c)* that $A \subseteq B$. Then

$$\mathrm{cl}(A) \subseteq \mathrm{cl}(A) \cup \mathrm{cl}(B) = \mathrm{cl}(A \cup B) = \mathrm{cl}(B)\,. \tag{3}$$

Now, if $\mathcal{Q} \subseteq \mathcal{F}$, then, for each $Q_0 \in \mathcal{Q}$, $\bigcap_{Q \in \mathcal{Q}} Q \subseteq Q_0$, and so, by (3),

$$\mathrm{cl}\Big(\bigcap_{Q \in \mathcal{Q}} Q\Big) \subseteq \mathrm{cl}(Q_0)\,.$$

This being so for each $Q_0 \in \mathcal{Q}$,

$$\mathrm{cl}\Big(\bigcap_{Q \in \mathcal{Q}} Q\Big) \subseteq \bigcap_{Q \in \mathcal{Q}} \mathrm{cl}(Q) = \bigcap_{Q \in \mathcal{Q}} Q$$

(recall that $\mathrm{cl}(Q) = Q$ for each $Q \in \mathcal{F}$, by the definition of $\mathcal{F}$). 2.1*(b)* gives the inclusion in the opposite direction; so $\bigcap_{Q \in \mathcal{Q}} Q \in \mathcal{F}$. This proves *(iii)*. □

Lemma 2.3. *Suppose that* $\mathcal{F}$ *is a class of subsets of* Ω *satisfying* 2.2(i)–(iii)*. Then there is a unique closure operation in* Ω *such that* $\mathcal{F}$ *is the class of its closed sets.*

Proof. Suppose $\mathcal{F}$ satisfies 2.2*(i)–(iii)*. Define, for any $A \in \mathcal{P}(\Omega)$,

$$\mathrm{cl}(A) := \bigcap_{\mathcal{F} \ni F \supseteq A} F\,. \tag{4}$$

Firstly, $A \subseteq \mathrm{cl}(A)$ (for it is the intersection of a class of sets all of which include A). Secondly, if $A \in \mathcal{F}$, A itself appears in the class whose intersection is taken, so that $A \supseteq \mathrm{cl}(A) \supseteq A$. This shows that $\mathrm{cl}(\emptyset) = \emptyset$, by 2.2*(i)*.

By 2.2*(iii),* $\mathrm{cl}(A) \in \mathcal{F}$, and, therefore, $\mathrm{cl}(\mathrm{cl}(A)) = \mathrm{cl}(A)$.

Finally, given $A, B \in \mathcal{P}(\Omega)$, $\mathrm{cl}(A) \subseteq \mathrm{cl}(A \cup B)$ directly from definition (the right-hand side is the intersection of a smaller class), and similarly for $\mathrm{cl}(B)$. Hence,

$$\mathrm{cl}(A) \cup \mathrm{cl}(B) \subseteq \mathrm{cl}(A \cup B)\,. \tag{5}$$

But $A \cup B \subseteq \mathrm{cl}(A) \cup \mathrm{cl}(B) \in \mathcal{F}$ by 2.2*(ii)*. Hence $\mathrm{cl}(A \cup B) \subseteq \mathrm{cl}(A) \cup \mathrm{cl}(B)$. With (5), this demonstrates 2.1*(c)*, and completes the proof that cl is a closure operation. We have already seen that $\mathrm{cl}(A) = A$ if and only if $A \in \mathcal{F}$.

That (4) defines the only possible closure operation giving rise to $\mathcal{F}$ as its class of closed sets is left as an *exercise*. □

The formula (4), with 2.2*(iii)*, means that $\mathrm{cl}(A)$ is the smallest member of $\mathcal{F}$ that includes A. This is the most memorable description of the procedure.

The properties 2.2*(i)–(iii)* recall 1.1*(a)–(c)*, with unions and intersections interchanged. It is therefore natural to make the following definitions.

Definition 2.4. *(a)* If Ω is a topological space with topology $\mathcal{G}$, a subset F of Ω is described as *closed* with respect to $\mathcal{G}$ if its complement $\Omega \setminus F$ belongs to $\mathcal{G}$. (In other words, a set is closed if its complement is open.) The class $\mathcal{F}$ of closed sets then gives rise to a closure operation defined by (4).

(b) If cl is a closure operation in the set Ω, Ω becomes a topological space if its topology $\mathcal{G}$ consists of the complements in Ω of sets closed with respect to cl.

A ritual remark here. Whichever definition you take as basic, *a set G is open if and only if its complement $\Omega \setminus G$ in Ω is closed*. Now, the complement is a set-theoretic construction: for a point x of Ω, $x \in G$ if and only if $x \notin \Omega \setminus G$. This definition is "at the level of points". However, it is not true (except in some special topologies, for instance the discrete topology) that G is open if and only if it is not closed — this would be the (usually false) statement "at the level of subsets" that $\mathcal{G} = \mathcal{P}(\Omega) \setminus \mathcal{F}$. That it is false in a familiar case like the standard topology on $\mathbb{R}$ is easily seen, for a set like $[0,1)$ is neither open nor closed.

In most of the interesting topological spaces, the class $\mathcal{G}$ of open sets (i.e. the topology) and the class $\mathcal{F}$ of closed sets are both "small" subclasses of $\mathcal{P}(\Omega)$, related to each other by $\mathcal{F} = \{\Omega \setminus G : G \in \mathcal{G}\}$ and, equivalently, $\mathcal{G} = \{\Omega \setminus F : F \in \mathcal{F}\}$. There is no *a priori* reason why a set should not be *both* open and closed, i.e. be a member of $\mathcal{F} \cap \mathcal{G}$. (Sets that are both open and closed do have some importance, for instance in the theory of Boolean algebras, and are occasionally called *clopen*, which is at least unambiguous. Notice that $\emptyset$ and Ω are *always* both open and closed.)

Example 2.5. Consider the space $\mathbb{C}^n$. A subset Z of $\mathbb{C}^n$ is a *variety* (strictly speaking a *complex affine variety*, the word "affine" meaning "in $\mathbb{C}^n$ rather than in complex projective space") if there is a finite set $\{p_1, p_2, \ldots, p_k\}$ of complex polynomials in n variables such that $Z = Z(p_1, p_2, \ldots, p_k) := \{x \in \mathbb{C}^n : p_i(x) = 0 \text{ for } 1 \le i \le k\}$. Thus $\mathbb{C}^n$ is itself a variety corresponding to the singleton of the zero polynomial, and $\emptyset$ is a variety corresponding to the singleton of any nonzero constant polynomial. The union of two varieties is a variety, for

$$\begin{aligned} &Z(p_1, \ldots, p_k) \cup Z(q_1, \ldots, q_l) = \\ &\qquad\qquad Z(p_1q_1, \ldots, p_1q_l, p_2q_1, \ldots, p_2q_l, \ldots, p_kq_1, \ldots, p_kq_l)\,. \end{aligned}$$

It is less obvious that the intersection of any class of varieties is still a variety. This follows from a non-trivial algebraic theorem that any ideal of the ring of complex polynomials in n unknowns has a finite ideal basis — one may formulate the concept of a "zero-set" Z in terms of an ideal of polynomials instead of a finite set thereof.

In any case, 2.2*(a)*, *(b)*, and *(c)* are all satisfied. The varieties are precisely the closed sets of a topology on $\mathbb{C}^n$. This topology is the *Zariski topology* on $\mathbb{C}^n$. It is very much coarser than the usual topology defined by the usual "Hermitian" metric. Indeed, when $n = 1$, the Zariski topology is just the cofinite topology.

There is nothing special about $\mathbb{C}$ here; any field would do.

Definition 2.6. A set E in the topological space Ω is *locally closed* if it is the intersection of an open set and a closed set. [This strange name will be clarified later.]

For example, the set $[0,1)$ in $\mathbb{R}$ is locally closed in the standard topology; for it is $[0,1]\cap(-1,1)$, or indeed $[0,100]\cap(-\frac{1}{100},1)$.

We have seen that a topology may be specified either directly, by defining the class of open sets itself, or by giving a base for it, or by giving a subbase, or by defining instead the class of closed sets (perhaps by a closure operation). There are still other methods, which are in a sense closer to the motivation for the theory. The notion of a topological space arose from the idea of “convergence” in the first place, and that could be expressed in terms of “neighbourhoods”.

Definition 2.7. Let Ω be a topological space, and $x\in\Omega$. A subset N of Ω is a *neighbourhood* of x in Ω if there is an open set U such that $x\in U\subseteq N$. The class consisting of all the neighbourhoods of x in Ω will be denoted by $\mathfrak{N}(x)$.

As often happens, this definition is not universally accepted. Some people insist that a neighbourhood must itself be open (a “neighbourhood of x” would then be just the same as an open set containing x). Authors occasionally write such things as “N is an open neighbourhood of x” either to clarify, by implication, that their convention is that of 2.7, or for verbal variety. More seriously, in a metric space, the word “neighbourhood” is quite often used to mean what I called in 312 a “ball”. Such a lack of complete agreement on terminology is common in mathematics (and not unknown in other subjects). The reasons are various, but you should be aware that if you open a mathematical book in the middle you may misunderstand the text because the author is using familiar words in somewhat different senses from those you have been accustomed to.

Lemma 2.8. *Let Ω be topological space, and let $\mathfrak{N}(x)$ denote the class of neighbourhoods of $x\in\Omega$. Then, for any $x\in\Omega$,*

(a) $(\forall N\in\mathfrak{N}(x))\quad x\in N$,

(b) $(\forall N\in\mathfrak{N}(x))\quad N\subseteq M\in\mathcal{P}(\Omega)\Longrightarrow M\in\mathfrak{N}(x)$,

(c) $(\forall N_1,N_2\in\mathfrak{N}(x))\quad N_1\cap N_2\in\mathfrak{N}(x)$,

(d) $(\forall N\in\mathfrak{N}(x))(\exists N_0\in\mathfrak{N}(x))\quad N_0\subseteq N\ \ \&\ \ (y\in N_0\Longrightarrow N\in\mathfrak{N}(y))$.

Proof. *(a)* and *(b)* are obvious; *(c)* follows from the fact that the intersection of open sets is open; and *(d)* is true if one takes N_0 to be the open set “U” of 2.7.

Of course $\mathfrak{N}(x)$ is specific to the space and the topology; I should really write $\mathfrak{N}(\Omega,\mathcal{G}_\Omega,x)$. But I shall usually leave it to the context to clarify what is meant.

Lemma 2.9. *Suppose that, to each element x of the set Ω, a nonempty class of subsets $\mathfrak{N}(x)$ is associated in such a way that the conditions* 2.8(a)–(d) *are satisfied. Then there is a unique topology $\mathcal{G}$ on Ω such that, for each x, $\mathfrak{N}(x)$ is the class of neighbourhoods of x in that topology.*

Proof. Let $\mathcal{G}:=\{U\in\mathcal{P}(\Omega):(\forall x\in U)\ U\in\mathfrak{N}(x)\}$. (In words: define a set in Ω to be open if it is a “neighbourhood” of each of its points).

Certainly $\emptyset\in\mathcal{G}$ (as it has no points, it is a neighbourhood of each!) For any $x\in\Omega$, we have required that $\mathfrak{N}(x)$ is nonempty; thus there is some $N\in\mathfrak{N}(x)$, and by 2.8*(b),* $N\subseteq\Omega$ and $\Omega\in\mathfrak{N}(x)$. By definition, $\Omega\in\mathcal{G}$. This proves 1.1*(a).*

Suppose $U_1, U_2 \in \mathcal{G}$. If $x \in U_1 \cap U_2$, then $U_1, U_2 \in \mathfrak{N}(x)$, and, by 2.8*(c)*, $U_1 \cap U_2 \in \mathfrak{N}(x)$. This proves 1.1*(b)*.

Similarly, suppose $\mathcal{U} \subseteq \mathcal{G}$, and $x \in \bigcup_{U \in \mathcal{U}} U$. There exists some $U_0 \in \mathcal{U}$ such that $x \in U_0$. But, as $U_0 \in \mathcal{G}$, $U_0 \in \mathfrak{N}(x)$. But $U_0 \subseteq \bigcup_{U \in \mathcal{U}} U$, so that, by 2.8*(b)*, $\bigcup_{U \in \mathcal{U}} U \in \mathfrak{N}(x)$. This proves 1.1*(c)*, and establishes that $\mathcal{G}$ is a topology in Ω.

Let M be a neighbourhood of $x \in \Omega$ in the topology $\mathcal{G}$. Thus, there is some $U \in \mathcal{G}$ with $x \in U \subseteq M$. But this implies, by the definition of $\mathcal{G}$, that $U \in \mathfrak{N}(x)$, and so, by 2.8*(b)*, that $M \in \mathfrak{N}(x)$. All the $\mathcal{G}$-neighbourhoods of x are in $\mathfrak{N}(x)$. So far 2.8*(d)* has been irrelevant.

Let $N \in \mathfrak{N}(x)$. Define

$$M := \{y \in N : N \in \mathfrak{N}(y)\}.$$

Take N_0 as in 2.8*(d)*. Then $N_0 \subseteq M$, so that, by 2.8*(b)*, $M \in \mathfrak{N}(x)$. On the other hand, if $y \in M$, then $N \in \mathfrak{N}(y)$, and 2.8*(d)* asserts the existence of $N_1 \in \mathfrak{N}(y)$ such that $N \in \mathfrak{N}(z)$ for every $z \in N_1$. Hence, $N_1 \subseteq M$, and $M \in \mathfrak{N}(y)$. This shows that M is a neighbourhood of every one of its points, i.e. that $M \in \mathcal{G}$; hence, N is a neighbourhood of x with respect to the topology $\mathcal{G}$. So $\mathfrak{N}(x)$ is precisely the class of all $\mathcal{G}$-neighbourhoods of x, for each $x \in \Omega$.

Finally, $\mathcal{G}$ must be the only topology with this property. Indeed, the open sets of such a topology must be neighbourhoods of every one of their elements. □

Clearly 2.8*(d)* is the “coherence” condition that ensures the various $\mathfrak{N}(x)$ fit together to define and be defined by a topology.

Since the topology may be fully described either by the classes $\mathfrak{N}(x)$ or by the closure operation, it is natural to use them to describe various constructions.

Definition 2.10. Let Ω be a topological space, and let A be a subset of Ω.

(a) A point $x \in \Omega$ is an *adherent point* [the name is not completely standard; Estate Khmaladze likes to call them *contact points*, for instance] of A if every neighbourhood of x meets A: that is, if $(\forall N \in \mathfrak{N}(x))\ N \cap A \neq \emptyset$.

(b) x is an *accumulation point* of A if every neighbourhood of x contains a point **other than** x: $(\forall N \in \mathfrak{N}(x))\ N \cap A \neq \{x\}$. [The name “accumulation point” is as neutral as I can manage. There are several very closely related concepts whose names depend on the author you are reading. Even worse, they are often equivalent in the more “natural” cases, so great care is necessary.] The set of accumulation points of A is called the *derived set* of A, and may be denoted by A'.

(c) A set $B \in \mathcal{P}(\Omega)$ is a *deleted* (or *punctured*) *neighbourhood* of $x \in \Omega$ if $x \notin B$ but $B \cup \{x\} \in \mathfrak{N}(x)$. [$B$ is a neighbourhood of x with x itself removed.]

(d) y is an *interior point* of A if A is a neighbourhood of y. The *interior* of A (denoted $\text{int}(A)$) is the set of interior points of A: $\text{int}(A) := \{x \in \Omega : A \in \mathfrak{N}(x)\}$.

(e) y is a *frontier point* of A if every neighbourhood of y meets both A and $\Omega \setminus A$. This means, of course, exactly the same as that $y \in \text{cl}(A) \cap \text{cl}(\Omega \setminus A)$. The set of frontier points of A is the *frontier* $\text{Fr}(A)$ of A:

$$\text{Fr}(A) := \text{cl}(A) \cap \text{cl}(\Omega \setminus A).$$

[The frontier is also called the *boundary* of A, and various other notations are used, such as $\dot{A}$ or ∂A. But both the words and the notations tend to be rather ambiguous, like these examples.]

Lemma 2.11. (a) *The closure of A is the set of adherent points of A.*

(b) $x \in A'$ *if and only if every punctured neighbourhood of x meets A.*

(c) *An adherent point of A is either an accumulation point of A or a point of A. Hence* $\mathrm{cl}(A) = A' \cup A$. *The union is not always* [or even "usually"] *disjoint.*

(d) $\mathrm{int}(A)$ *is the largest open set included in A:* $\mathrm{int}(A) = \bigcup_{A \supseteq U \in \mathcal{G}} U$.

(e) $\mathrm{int}(A) = \Omega \setminus \mathrm{cl}(\Omega \setminus A)$. [If we write A^c for the complement of A, we can more succinctly state: $\mathrm{int}(A^c) = (\mathrm{cl}(A))^c$, and various equivalent formulations.]

Proof. I prove *(a)*, leaving the rest as an exercise. Suppose $x \notin \mathrm{cl}(A)$, which is closed. Then $\Omega \setminus \mathrm{cl}(A)$ is a neighbourhood of x that does not meet $A \subseteq \mathrm{cl}(A)$.

On the other hand, if $N \in \mathfrak{N}(x)$ and $N \cap A = \emptyset$, there is an open set U with $x \in U \subseteq N$, and *a fortiori* $U \cap A = \emptyset$; therefore, $\Omega \setminus U$ is closed and includes A, which implies that $x \notin \Omega \setminus U \supseteq \mathrm{cl}(A)$. □

Remark 2.12. We saw in 2.9 that a set A in Ω is open if and only if A is a neighbourhood of each of its points; this says precisely that $\mathrm{int}(A) = A$. Similarly, A is closed if and only every adherent point of A belongs to B. These facts are often useful; they are perhaps closer to our intuition about the concepts than the definitions we gave.

As with topologies, we can specify $\mathfrak{N}(x)$ more economically.

Definition 2.13. *(a)* A *base of neighbourhoods* at a point x of the topological space Ω is a subclass $\mathfrak{B}(x)$ of $\mathfrak{N}(x)$ such that, for any $N \in \mathfrak{N}(x)$, there is some $B \in \mathfrak{B}(x)$ for which $B \subseteq N$.

(b) A *subbase of neighbourhoods* at $x \in \Omega$ is a subclass $\mathcal{S}(x)$ of $\mathfrak{N}(x)$ such that the class of finite intersections of members of $\mathcal{S}(x)$ (that is, of sets of the form $N_1 \cap N_2 \cap \ldots \cap N_k$, where $k \in \mathbb{N}$ and $N_1, N_2, \ldots, N_k \in \mathcal{S}(x)$) is a base of neighbourhoods at x.

(c) Ω is said to satisfy the *first axiom of countability* (or, vulgarly, to be *first countable*) if every point has a countable base of neighbourhoods. [In fact it suffices to know it has a countable subbase of neighbourhoods, as we show next.]

Lemma 2.14. (a) *If in a topological space Ω the point x has a countable subbase of neighbourhoods, it has a countable base of neighbourhoods.*

(b) *If $x \in O$ has a countable base of neighbourhoods, then it has a base of neighbourhoods that may be listed as a decreasing sequence of sets.*

Proof. *(a)* Let $\mathcal{S}(x)$ be a countable subbase of neighbourhoods of x, and enumerate it as $\{S_1, S_2, S_3, \ldots\}$. The class $\mathcal{J}$ of all *finite* subsets of $\mathbb{N}$ is countable (why?), and the mapping $\mathcal{J} \ni J \mapsto \bigcap_{j \in J} S_j$ is onto the class of all finite intersections of members of $S(x)$. So the class of all the finite intersections of members of $\mathcal{S}(x)$ is a countable base of neighbourhoods of x.

(b) By hypothesis, there is a base of neighbourhoods $\{U_n : n \in \mathbb{N}\}$ for x. For each $n \in \mathbb{N}$, let $V_n := U_1 \cap U_2 \cap \cdots \cap U_n$. It is clear that $V_1 \supseteq V_2 \supseteq \cdots$, that each V_n is a neighbourhood of x, and that $\{V_n : n \in \mathbb{N}\}$ is a base of neighbourhoods for x, since $V_n \subseteq U_n$ for each n. □

It is clear (why?) that a second countable space is first countable. The converse is not necessarily true.

Lemma 2.15. *A metric space is first countable.*

Proof. The class $\{B(x;1/n) : n \in \mathbb{N}\}$ is a countable base for $\mathfrak{N}(x)$. $\square$

We shall soon see that there are useful metric spaces that are *not* second countable.

We have now met a number of ways of describing closures and interiors of a set A, and each of them could be taken as the definition. It would, for instance, be quite possible to define a topological space as a set Ω with an "interior" operation in $\mathcal{P}(\Omega)$ having suitable properties (complementary to the "closure" properties of 2.1).

Definition 2.16. Let A and B be subsets of the topological space Ω. We say that A is *dense* in B if $B \subseteq \mathrm{cl}_\Omega(A)$. [Again, this terminology is logically superfluous, but quite common. Usually it is applied when $A \subseteq B$.]

§3. Continuity.

In this section Ω, Ψ, Φ will be topological spaces, with the topologies $\mathcal{G}_\Omega, \mathcal{G}_\Psi, \mathcal{G}_\Phi$ respectively. I start by labouring a point that I seem to make in all my courses.

Definition 3.1. Suppose that $f : \Omega \longrightarrow \Psi$ is a mapping and $B \in \mathcal{P}(\Psi)$. Define

$$f^{-1}(B) := \{y \in \Omega : f(y) \in B\}. \tag{6}$$

Thus, $f^{-1}(B)$ is a subset (possibly null) of Ω, called the *inverse image of B under f*.

If $A \in \mathcal{P}(\Omega)$, we define

$$f(A) := \{f(x) : x \in A\}, \tag{7}$$

which is a subset of Ψ called the *image of A under f*.

These notations, although they are quite standard, are genuinely confusing, because (7) might lead one to interpret (6) as the image of B under an "inverse mapping" f^{-1} from Ψ to Ω. But the formula (6) *always* makes sense, whether or not f has an inverse mapping.

Recall that $f : \Omega \longrightarrow \Psi$ has an inverse mapping from Ψ to Ω if and only if it is a bijection. In that rather unusual case, the inverse mapping would commonly be denoted f^{-1}, and then the expression $f^{-1}(B)$ could be interpreted in two ways: as the image of B under the mapping f^{-1}, as at (7), or according to (6), as the inverse image of B under f. However, both interpretations would yield the same subset of Ω. That being so, there is never, in fact, any *practical* confusion in using the notation (6), provided one bears in mind that it does not imply the existence of a mapping inverse to f.

To make the *theoretical* confusion worse, the formula (6) establishes a mapping

$$B \longmapsto f^{-1}(B) : \mathcal{P}(\Psi) \longrightarrow \mathcal{P}(\Omega),$$

which it would be absurd not to call f^{-1}. Thus the notation f^{-1} does in all cases denote a legitimate mapping, but it is between the power classes rather than between the original sets. Similarly, (7) defines a mapping

$$A \longmapsto f(A) : \mathcal{P}(\Omega) \longrightarrow \mathcal{P}(\Psi)$$

which might simply be called f, although it is not the same as the "f" we started with.

Lemma 3.2. *Suppose* f, Ω, Ψ *are as above; and let* $\{A_\gamma : \gamma \in \Gamma\}$ *be an indexed family of subsets of* Ω, $\{B_\gamma : \gamma \in \Gamma\}$ *be an indexed family of subsets of* Ψ, *and let* $B, B' \in \mathcal{P}(\Psi)$. *Then*

$$f^{-1}\Big(\bigcup_{\gamma\in\Gamma} B_\gamma\Big) = \bigcup_{\gamma\in\Gamma} f^{-1}(B_\gamma)\,, \qquad f^{-1}\Big(\bigcap_{\gamma\in\Gamma} B_\gamma\Big) = \bigcap_{\gamma\in\Gamma} f^{-1}(B_\gamma)\,,$$
$$f^{-1}(B \setminus B') = f^{-1}(B) \setminus f^{-1}(B')\,, \qquad f^{-1}(\Psi) = \Omega\,, \qquad f^{-1}(\emptyset) = \emptyset\,,$$
$$f\Big(\bigcup_{\gamma\in\Gamma} A_\gamma\Big) = \bigcup_{\gamma\in\Gamma} f(A_\gamma)\,, \qquad f\Big(\bigcap_{\gamma\in\Gamma} B_\gamma\Big) \subseteq \bigcap_{\gamma\in\Gamma} f(B_\gamma)\,. \qquad \square$$

It is also obvious that $f(\emptyset) = \emptyset$, but $f(\Omega) = \Psi$ only if f is onto (by definition). In short, $f^{-1} : \mathcal{P}(\Psi) \longrightarrow \mathcal{P}(\Omega)$ respects *all* the set-theoretic operations, including the "unary" operations $\emptyset$ and "whole space", but $f : \mathcal{P}(\Omega) \longrightarrow \mathcal{P}(\Psi)$ need not.

Notice that, if $A, A' \in \mathcal{P}(\Omega)$ and $A \subseteq A'$, then $f(A) \subseteq f(A')$; similarly, if $B \subseteq B'$ above, then $f^{-1}(B) \subseteq f^{-1}(B')$. These statements follow from 3.2, but are almost obvious anyway.

Lemma 3.3. *Suppose that* $f : \Omega \longrightarrow \Psi$ *and* $g : \Psi \longrightarrow \Phi$, $A \in \mathcal{P}(\Omega)$, *and* $B \in \mathcal{P}(\Phi)$. *Then* $(g \circ f)(A) = g(f(A))$ *and* $(g \circ f)^{-1}(B) = f^{-1}(g^{-1}(B))$. $\square$

I shall normally write gf instead of $g \circ f$, unless I need to avoid ambiguity.

Definition 3.4. Let $f : \Omega \longrightarrow \Psi$ be any mapping. f is *continuous at* $x \in \Omega$ if, for any neighbourhood N of $f(x)$ in Ψ (in the topology $\mathcal{G}_\Psi$), the set

$$f^{-1}(N) := \{y \in \Omega : f(y) \in N\} \tag{8}$$

is a neighbourhood of x in Ω (in the topology $\mathcal{G}_\Omega$).

We say that f is *continuous* (without further qualification) if, for every open set U in Ψ, $f^{-1}(U)$ is an open set in Ω.

Remark 3.5. The (ϵ, δ) definition of continuity, at a point, of a function between metric spaces is easily seen to be equivalent in those circumstances to the above.

Lemma 3.6. f *is continuous at* x *if and only if, for any closed set* C *in* Ψ *such that* $f(x) \notin C$, $x \notin \mathrm{cl}_\Omega(f^{-1}(C))$. $\square$

Lemma 3.7. *Let* $\mathcal{S}$ *be a subbase for* $\mathcal{G}_\Psi$. *A mapping* $f : \Omega \longrightarrow \Psi$ *is continuous if and only if either of the following equivalent conditions is satisfied.*

(a) *For every closed subset* C *of* Ψ, *the inverse image* $f^{-1}(C)$ *is closed in* Ω.

(b) *For every* $U \in \mathcal{S}$, $f^{-1}(U) \in \mathcal{G}_\Omega$.

Proof. *(a)* If C is closed, $\Psi \setminus C$ is open, and, by 3.2, $f^{-1}(\Psi \setminus C) = \Omega \setminus f^{-1}(C)$ is open if and only if $f^{-1}(C)$ is closed; and so on. Similarly for *(b)*. $\square$

Lemma 3.8. *$f : \Omega \longrightarrow \Psi$ is continuous if and only if it is continuous at each point of Ω.*

Proof. Suppose that f is continuous at each point of Ω, and let U be an open set of Ψ. If $x \in f^{-1}(U)$, then $f(x) \in U$, so $U \in \mathfrak{N}(f(x))$ and (as f is continuous at x) $f^{-1}(U)$ is a neighbourhood of x. But this shows that every element of $f^{-1}(U)$ is an interior point, so that $f^{-1}(U)$ is open in Ω.

Conversely, if f is continuous and $x \in \Omega$ and $N \in \mathfrak{N}(f(x))$, there is some V, open in Ψ, such that $f(x) \in V \subseteq N$. Then $x \in f^{-1}(V) \subseteq f^{-1}(N)$, where $f^{-1}(V)$ is open (as f is continuous). This means that $f^{-1}(N) \in \mathfrak{N}(x)$. □

Lemma 3.9. (a) *Let $f : \Omega \longrightarrow \Psi$ and $g : \Psi \longrightarrow \Phi$. If $x \in \Omega$, f is continuous at x, and g is continuous at $f(x)$, then gf is continuous at x. If f and g are continuous, so is gf.*

(b) *The identity map of any topological space is continuous.*

(c) *Any constant map (i.e. one whose image is a singleton) is continuous.*

Proof. *(a)* follows almost instantly from 3.3. For *(b)*, recall that the identity map $\mathbf{1}_\Omega$ of Ω is defined by setting $\mathbf{1}_\Omega(x) := x$ for all $x \in \Omega$. Thus, for any open set U in Ω, $\mathbf{1}_\Omega^{-1}(U) = U$, which is also open. This proves that $\mathbf{1}_\Omega$ is continuous.

For *(c)*, suppose that $f(\Omega) = \{c\}$, or equivalently that $f(x) = c$ for all $x \in \Omega$. Let $U \in \mathcal{G}_\Psi$. Then, if $c \in U$, $f^{-1}(U) = \Omega$, and, if $c \notin U$, $f^{-1}(U) = \emptyset$. In either case, $f^{-1}(U)$ is open in Ω. □

Those of you who are taking 439 will recognize that *(a)* and *(b)* say precisely that topological spaces and continuous maps between them form a *category*, since composition of mappings is always associative. As I may possibly have a little more to say about categories later, let me just comment that:

Definition 3.10. A *category* may be described as follows. There is a class of *objects*. For any two objects A and B of the category, there is given a set $Mor(A, B)$ of *morphisms* "from A to B". If A, B, C are objects of the category, there is a mapping

$$Mor(A, B) \times Mor(B, C) \longrightarrow Mor(A, C) : (f, g) \mapsto gf$$

(called "composition" of morphisms), which is associative: if $f \in Mor(A, B)$, $g \in Mor(B, C)$, and $h \in Mor(C, D)$, then $(hg)f = h(gf)$. Further, for each object B, there is an "identity" morphism $\mathbf{1}_B \in Mor(B, B)$ which has the properties that $\mathbf{1}_B f = f$ for any $f \in Mor(A, B)$ and $g\mathbf{1}_B = g$ for any $g \in Mor(B, C)$.

I have formulated this definition a little imprecisely, since set-theoretical niceties arise if one tries to be more careful. Apart from topological spaces and continuous maps, other examples of categories are: (all) real vector spaces and linear maps between them; abelian groups and (group) homomorphisms between them; (all) groups and (group) homomorphisms between them; rings and ring homomorphisms; sets and mappings between them; $\mathbb{R}$ and increasing functions $\mathbb{R} \longrightarrow \mathbb{R}$; two complex vector spaces E and F, the linear maps $E \longrightarrow F$, and the identity mappings of E and of F. It is possible to have only one object; the sets of morphisms may also be small (and even sometimes empty). Only identity morphisms, one for each object, *have* to exist.

Lemma 3.11. *Suppose that $\mathcal{G}_1$ and $\mathcal{G}_2$ are topologies on the same space Ω. Then $\mathcal{G}_1$ is finer than $\mathcal{G}_2$, $\mathcal{G}_1 \supseteq \mathcal{G}_2$, if and only if the identity map $(\Omega, \mathcal{G}_1) \longrightarrow (\Omega, \mathcal{G}_2)$ is continuous.* $\square$

Definition 3.12. Let Ω and Ψ be topological spaces. A mapping $f : \Omega \longrightarrow \Psi$ is a *homeomorphism* (well, really a homœomorphism, but the American spelling is generally accepted nowadays) if it is continuous and a bijection, and its inverse map is also continuous. That is: f is a homeomorphism if is continuous and there exists a continuous mapping $g : \Psi \longrightarrow \Omega$ such that $fg = \mathbf{1}_\Psi$ and $gf = \mathbf{1}_\Omega$.

[That is: homeomorphisms are the *isomorphisms* in the category of topological spaces and continuous maps.]

A continuous mapping, as defined above, is one for which the inverse image of an open [or closed] set is open [or closed]. There are also classes of mappings defined by the behaviour of (direct) images, and their definitions may as well be given here:

Definition 3.13. A mapping $f : \Omega \longrightarrow \Psi$ between topological spaces is

(a) *open* if, for every open set U in Ω, $f(U)$ is an open set in Ψ; and

(b) *closed* if, for every closed set C in Ω, $f(C)$ is a closed set in Ψ.

Clearly these conditions have no general relation to continuity. But a bijection will be open (and closed) if and only if its inverse is continuous.

Lemma 3.14. *If the mapping $f : \Omega \longrightarrow \Psi$ between topological spaces is continuous, then, for any $A \in \mathcal{P}(\Omega)$, $f(\mathrm{cl}_\Omega(A)) \subseteq \mathrm{cl}_\Psi(f(A))$. f is closed if the inclusion is always an equality.* $\square$

§4. The axiom of choice.

I do not wish to spend much time on the axiom of choice as such, because it is presumably fully treated in the logic courses. But the "pure" axiom is often not the most convenient version to use. There are other statements that are often more useful and in a sense equivalent.

The "pure" Axiom of Choice is the assertion that, for any set $\mathcal{D}$ whose members are non-empty sets, there is a function $f : \mathcal{D} \longrightarrow \bigcup_{C \in \mathcal{D}} C$ such that $f(C) \in C$ for each $C \in \mathcal{D}$. (In other words, the "choice function" f "chooses" one element from each of the sets $C \in \mathcal{D}$. Notice that $\mathcal{D}$ is the domain of the function f; the *arguments* of f are the members of $\mathcal{D}$, and the values of f are elements of members of $\mathcal{D}$.)

If $\mathcal{D}$ is a singleton, there is certainly a "choice function", for its only member C (being nonnull) has an element c (that being what it means to say $C \neq \emptyset$) and then there is a choice function whose only value $f(C)$ is c. In any of the standard axiom systems for formal set theory (like NBG or ZF) one can prove similarly that a choice function exists if $\mathcal{D}$ is finite. Whatever the logical problems involved in formalizing the ideas, I think few people would hesitate to assume the existence of a choice function when $\mathcal{D}$ is countably infinite (perhaps we feel that the only difficulty in making an infinite sequence of choices is the shortness of life). But the historical prominence of the Axiom of Choice is because it asserts the existence of a choice function even in cases where $\mathcal{D}$ is far too big to conceive of any "practical" way to construct one.

At the merely intuitive level, I find the Axiom unexceptionable, and perhaps no-one would have worried very much about it had not Russell constructed his paradox and thereby demonstrated the pitfalls of naive set theory. We do know, now, that the axiom is consistent (Gödel), as is its negation (Mendelson), with NBG or ZF.

There is a large number of statements which were proved quite early to be equivalent to the Axiom of Choice, in the sense that they may all be derived from each other if you assume the standard axioms of set theory. I shall not give proofs; you can find a concise discussion in the appendix to Kelley's book *General Topology* (it is unimportant that he begins from a stronger axiom) and in many other places, and, in truth, we do not need more than a smattering of information. The standard joke is that "mathematics is independent of its foundations"; provided our intuition approves and no paradoxes arise, we should not worry too much about the details.

A. The *Multiplicative Axiom* is the assertion that, if $\{C_\alpha : \alpha \in A\}$ is an indexed class of non-empty sets, the product $\prod_{\alpha\in A} C_\alpha$ (or $\bigtimes_{\alpha\in A} C_\alpha$) is also non-empty.

What is at issue here is really the definition of the Cartesian product of a general indexed family, not just of finitely many sets $C_1, C_2, \ldots, C_k$. The usual convention is that $\prod_{\alpha\in A} C_\alpha$ is *by definition* the set of "choice functions" $A \longrightarrow \bigcup_{\alpha\in A} C$ for which $f(\alpha) \in C_\alpha$ for each $\alpha \in A$. [Thus, $\prod_{\alpha=1,2} C_\alpha$ is the set of functions $f : \{1,2\} \longrightarrow C_1 \cup C_2$ for which $f(1) \in C_1$ and $f(2) \in C_2$. It is clear that this is, as it were, "functionally equivalent" to the definition of $C_1 \times C_2$ as a set of ordered pairs, although it is *not* the same.] This being so, the Multiplicative Axiom as I have stated it is *almost* a rephrasing of the Axiom of Choice as I have stated it, and they are very nearly trivially equivalent. [There is a difference; what is it?]

B. The *Well-Ordering Principle* (also known as *Zermelo's Axiom*, or indeed *Zermelo's Theorem*) states that any set may be well-ordered.

A *partial order* on a set S is a binary relation $\leq$ on S which is reflexive ($a \leq a$) and transitive ($a \leq b$ & $b \leq c \Longrightarrow a \leq c$), such that $a \leq b$ & $b \leq a \Longrightarrow a = b$. (The last condition excludes trivial examples.) A *well-ordering* on S is a partial order such that every subset of S has a least element with respect to $\leq$.

The equivalence of the Axiom of Choice and the Well-Ordering Principle is well-known; in essence it goes back to Zermelo (1908), and may be found in many places — a sketch is in Kelley's book, for instance. The WOP is interesting because it allows one to carry out arguments by so-called transfinite induction, which extends ordinary induction, and so raises all the problems of transfinite arithmetic.

C. *Zorn's Lemma* is perhaps the most generally useful of these statements, and is so often used that people speak jovially of Zornifying.

Let $(S, \leq)$ be any partially ordered set. A *chain* in S is a subset T that is totally ordered by $\leq$ (that is, if $x, y \in T$, then either $x \leq y$ or $y \leq x$). An *upper bound* for T is an element $b \in S$ such that, for any $x \in T$, $x \leq b$. A *maximal element* of S is an element $m \in S$ such that $m \leq y \in S \Longrightarrow y = m$. (Notice that it need not be an upper bound for S. For instance, let S be $\{\{1\},\{1,2\},\{1,3\}\}$, with $\subseteq$ as $\leq$. Then $\{1,2\}$ and $\{1,3\}$ are both maximal in S, but neither is greater than the other.) The partial order $\leq$ on S is *inductive* if every chain in S has an upper bound. Zorn's Lemma then says that

an inductive partial order on a non-empty set has a maximal element.

The statement is superficially plausible. If there were no maximal element, then any $c \in S$ must allow a greater element, and one could construct a "sequence" of elements of S that is a chain and must have an upper bound. If a large enough chain exists, its upper bound should be maximal in S. This is where the Axiom of Choice is needed (to construct such a chain). Again, Zorn's Lemma is equivalent to the Axiom, and proofs in both directions may be found in many places. At first sight the Lemma seems over-elaborate, but it can be applied in very diverse situations by suitable choices of the partially ordered set.

It should be emphasized that S has to be a *set* for Zorn's Lemma to work. (Otherwise one can easily derive paradoxes such as Cantor's or Burali-Forti's.)

There are at least three other statements equivalent to the Axiom of Choice (the Hausdorff maximal principle, Kuratowski's maximal principle, and the minimal principle) that are very similar to Zorn's Lemma, and are in fact special cases of it. There is also another form of the Axiom of Choice (the Teichmüller-Tukey Lemma) that can easily be deduced from Zorn's Lemma. So, for many purposes, Zorn's lemma is the appropriate form of the Axiom of Choice.

It cannot be too strongly emphasized that the Axiom of Choice or its equivalents are only likely to be significant in "abstract and general" situations. For instance, we need them to prove that *every* vector space has a basis in the algebraic sense, because there are vector spaces so enormous that describing a basis explicitly would be impossible. In "practical" mathematics, one is unlikely to need more than the countable Axiom, which, as I said, seems intuitively acceptable. These days, however, we tend to apply the Axiom freely, since no (new) contradictions can result.

§5. New topological spaces from old.

Whenever mathematicians define a new concept, they ask how it can be manipulated. For topological spaces, there are many possible ways of obtaining new ones from old.

Definition 5.1. Let A be a subset of the topological space Ω (whose topology is $\mathcal{G}_\Omega$). The *subspace topology* (sometimes called the relative topology) on A *induced* from the topology on Ω is

$$\mathcal{G}_A := \{A \cap U : U \in \mathcal{G}_\Omega\}\,. \tag{9}$$

In words: a set is open in the subspace topology on A if it is the intersection with A of an open set of Ω. Of course we must check

Lemma 5.2. *$\mathcal{G}_A$, as defined by* (9), *is a topology on A.* □

Remark 5.3. If $A \in \mathcal{G}_\Omega$, then $\mathcal{G}_A = \{U \in \mathcal{G}_\Omega : U \subseteq A\}$ (and conversely); that is, the "relatively open" subsets of an open subset A of Ω are exactly the same as the subsets of A that are open in Ω. Similarly, if A is closed in Ω, the subsets of A that are relatively closed in A are precisely the subsets of A that are closed in Ω.

In cases like this, it is reassuring, especially from the viewpoint of category theory, if the new construct can be "characterized" by a "universal property"; or if it is "unique" in some suitable class. Recall that the *inclusion map* $i_A : A \longrightarrow \Omega$ is defined by setting $i_A(x) := x$ for each $x \in A$.

Lemma 5.4. (a) $\mathcal{G}_A$ *is the coarsest topology on* A *such that the inclusion map* i_A *is continuous (when* Ω *has the given topology* $\mathcal{G}_\Omega$*). That is: if* $\mathcal{G}$ *is a topology on* A *with the property that* $i_A : (A, \mathcal{G}) \longrightarrow (\Omega, \mathcal{G}_\Omega)$ *is continuous, then* $\mathcal{G} \supseteq \mathcal{G}_A$.

(b) *A mapping* $f : (\Psi, \mathcal{G}_\Psi) \longrightarrow (A, \mathcal{G}_A)$, *where* $(\Psi, \mathcal{G}_\Psi)$ *is any topological space, is continuous if and only if* $i_A f : (\Psi, \mathcal{G}_\Psi) \longrightarrow (\Omega, \mathcal{G}_\Omega)$ *is continuous.*

(c) $\mathcal{G}_A$ *is the only topology on* A *for which the statement* (b) *is true.*

Proof. *(a)* Given $U \in \mathcal{G}_\Omega$, $i_A^{-1}(U) = U \cap A$. Hence, if $i_A : (A, \mathcal{G}) \longrightarrow (\Omega, \mathcal{G}_\Omega)$ is continuous, $U \cap A \in \mathcal{G}$ for each such U, and, from (9), $\mathcal{G}_A \subseteq \mathcal{G}$. Conversely, if $\mathcal{G}_A \subseteq \mathcal{G}$, $i_A^{-1}(U) = U \cap A \in \mathcal{G}_A$ for each $U \in \mathcal{G}_\Omega$, so that i_A is continuous.

(b) $i_A f$ is continuous if and only if $(i_A f)^{-1}(U) \in \mathcal{G}_\Psi$ for any $U \in \mathcal{G}_\Omega$. But by 3.3, $(i_A f)^{-1}(U) = f^{-1}(i_A^{-1}(U))$, and every open set of $\mathcal{G}_A$ is $i_A^{-1}(U) = U \cap A$ for some $U \in \mathcal{G}_\Omega$. Thus the conditions that $i_A f$ be continuous and that f be continuous as a map into $(A, \mathcal{G}_A)$ are the same.

(c) Let $\mathcal{G}$ be a topology on A such that *(b)* is true (if $\mathcal{G}$ is substituted for $\mathcal{G}_A$). Then, taking $f := \mathbf{1}_A$, which is certainly continuous $(A, \mathcal{G}) \longrightarrow (A, \mathcal{G})$, we must have $i_A f = i_A$ is continuous, so that $\mathcal{G} \supseteq \mathcal{G}_A$ by *(a)*.

However, $i_A \mathbf{1}_A : (A, \mathcal{G}_A) \longrightarrow (\Omega, \mathcal{G}_\Omega)$ is also continuous by *(a)*, so that, from our assumption of *(b)*, $\mathbf{1}_A : (A, \mathcal{G}_A) \longrightarrow (A, \mathcal{G})$ is continuous. By 3.11, $\mathcal{G}_A \supseteq \mathcal{G}$.

The two inclusions together show that $\mathcal{G}_A = \mathcal{G}$. ☐

One often says simply "A is a (topological) *subspace* of the topological space Ω", meaning both that A is a subset of Ω and that it is understood to be furnished with the subspace topology induced from the topology of Ω.

Lemma 5.5. *Let* Ψ *be a subspace of the topological space* Ω*, and* $A \subseteq \Psi$. *Then*

$$\mathrm{cl}_\Psi(A) = \Psi \cap \mathrm{cl}_\Omega(A).$$ ☐

Definition 5.6. Let (Ω, d) be a metric space, and $A \subseteq \Omega$. Define a metric d_A on A to be the restriction of the metric d on Ω, $d_A := d|\, A \times A$; that is

$$(\forall x, y \in A) \quad d_A(x, y) := d(x, y).$$

Then (A, d_A) is described as a *metric subspace* of (Ω, d). (Of course one rarely bothers with the distinct notation d_A.) One then has

Lemma 5.7. *The topology defined by the subspace metric* d_A *is the subspace topology induced from the topology defined by the metric* d *on* Ω. ☐

Of course, if the Lemma had not been true, we should have considered that the definitions were inappropriate.

The general idea of the subspace topology can be extended.

Definition 5.8. Let $f : \Omega \longrightarrow \Psi$ be any mapping whatever, and let $\mathcal{H}_\Psi$ be a topology on Ψ. Define (the notation is just for temporary convenience)

$$\mathcal{G}^f := \{f^{-1}(U) : U \in \mathcal{H}_\Psi\} .$$

Then it follows trivially from 3.2 that $\mathcal{G}^f$ is a topology on Ω, the *topology on* Ω *induced by* f (from $\mathcal{H}_\Psi$).

Lemma 5.9. (a) *If* Ψ *has the topology* $\mathcal{H}_\Psi$*, then* $\mathcal{G}^f$ *is the coarsest in the class of topologies* $\mathcal{G}$ *on* Ω *such that* $f : (\Omega, \mathcal{G}) \longrightarrow (\Psi, \mathcal{H}_\Psi)$ *is continuous.*

(b) *Let* $(\Phi, \mathcal{G}_\Phi)$ *be any topological space. A mapping* $g : \Phi \longrightarrow (\Omega, \mathcal{G}^f)$ *is continuous if and only if* $fg : (\Phi, \mathcal{G}_\Phi) \longrightarrow (\Psi, \mathcal{H}_\Psi)$ *is continuous.*

(c) $\mathcal{G}^f$ *is the only topology on* Ω *for which* (b) *is true for all mappings* g. □

Lemma 5.10. *Suppose that* f *is a surjection. Then* $f : (\Omega, \mathcal{G}^f) \longrightarrow (\Psi, \mathcal{H}_\Psi)$ *is not only continuous, but also open.* □

However, there is no essential reason to restrict attention to a single mapping.

Lemma 5.11. *Let* $\mathfrak{G}$ *be a class of topologies on the set* Ω*. There is a coarsest topology* $\mathcal{G}^0$ *which is finer than all the topologies in* $\mathfrak{G}$*.*

Proof. If $\mathfrak{G} \neq \emptyset$, it is only necessary to take $\bigcup_{\mathcal{G}\in\mathfrak{G}} \mathcal{G}$ as a subbase for $\mathcal{G}^0$ (the condition of 1.16 is clearly satisfied). If $\mathfrak{G} = \emptyset$, the indiscrete topology is the coarsest possible. □

Hence, in the partial order in the class of all topologies on Ω ($\mathcal{G}_1 \leq \mathcal{G}_2$ means $\mathcal{G}_2$ is finer than $\mathcal{G}_1$, $\mathcal{G}_1 \subseteq \mathcal{G}_2$), *any subclass whatever* has a least upper bound in the class. [The partial order is "complete". This is at first glance better than Dedekind's axiom, which deals with non-empty subsets that are bounded above (or below). But in this case the class of *all* topologies anyway has an upper bound, the discrete topology, and a lower bound, the indiscrete topology.] It follows, as I proved in 312, that any subclass also has a greatest lower bound — but the "abstract nonsense" proof I gave then did not tell us what the g.l.b. is.

Lemma 5.12. *Let* $\mathfrak{G}$ *be a class of topologies on the set* Ω*. There is a finest topology* $\mathcal{G}_0$ *which is coarser than all the topologies in* $\mathfrak{G}$*.*

Proof. Take $\mathcal{G}_0 := \bigcap_{\mathcal{G}\in\mathfrak{G}} \mathcal{G}$. It is easily checked that this is a topology, and it is evidently the finest possible topology that is coarser than all the $\mathcal{G} \in \mathfrak{G}$. □

Corollary 5.13. *Let* Ω *be any set, and let* $\mathcal{F}$ *be any family of mappings* $f : \Omega \longrightarrow \Psi_f$*, where each of the codomains* Ψ_f *is a topological space with topology* $\mathcal{H}_f$*. There is a coarsest topology* $\mathcal{G}$ *on* Ω *such that each mapping in* $\mathcal{F}$ *is continuous.*

Proof. Each $f \in \mathcal{F}$ determines a topology $\mathcal{G}^f$ on Ω (by 5.8), and 5.11 constructs the coarsest topology finer than all the $\mathcal{G}^f$, which is the desired topology $\mathcal{G}$. $\square$

This is of course "abstract nonsense" again. The topology $\mathcal{G}$ can be described "explicitly" by specifying a base. It will consist of all the sets

$$f_1^{-1}(U_1) \cap f_2^{-1}(U_2) \cap \cdots \cap f_k^{-1}(U_k)$$

where k is any natural number and, for each i, $1 \leq i \leq k$, $U_i \in \mathcal{H}_{f_i}$.

An analyst would describe the topology $\mathcal{G}$ on Ω as the *weak topology* on Ω defined by the family $\mathcal{F}$ of mappings. If Ω already has a topology $\mathcal{H}$ and the mappings in $\mathcal{F}$ are continuous with respect to $\mathcal{H}$, then the weak topology $\mathcal{G}$ is coarser than $\mathcal{H}$.

A striking example would be if $\mathcal{F}$ consisted of all continuous real-valued functions on the topological space Ω. The resulting weak topology will have exactly the same real-valued functions as the original topology $\mathcal{H}$, but it may be much coarser if $\mathcal{H}$ only allows rather few continuous functions into $\mathbb{R}$.

Example 5.14. Let Ω be an infinite set; give it the cofinite topology $\mathcal{H}$ (see 1.4*(a)*). Suppose $f : \Omega \longrightarrow \mathbb{R}$ is continuous.

For any $\alpha \in \mathbb{R}$, $f^{-1}((-\infty, \alpha)) \in \mathcal{H}$, so must be empty or cofinite; likewise $f^{-1}((\alpha, \infty))$. As $f^{-1}((-\infty, \alpha)) \cap f^{-1}((\alpha, \infty)) = \emptyset$, they cannot both be cofinite.

Consider $A := \{\alpha \in \mathbb{R} : f^{-1}((-\infty, \alpha)) = \emptyset\}$. Clearly $\beta < \alpha \in A \Rightarrow \beta \in A$; and $A \neq \mathbb{R}$, since $\emptyset \neq \Omega = f^{-1}(\mathbb{R}) = \bigcup_{\alpha \in \mathbb{R}} f^{-1}((-\infty, \alpha))$. But also $A \neq \emptyset$; for, if $A = \emptyset$, then $f^{-1}((-\infty, \alpha))$ would be cofinite and $f^{-1}((\alpha, \infty))$ would have to be empty for all α, which would lead to the same absurd conclusion as before that $f^{-1}(\mathbb{R}) = \bigcup_{\alpha \in \mathbb{R}} f^{-1}((\alpha, \infty)) = \emptyset$.

Hence $A = (-\infty, \beta)$ or $A = (-\infty, \beta]$ for some $\beta := \sup A \in \mathbb{R}$. [This is, I hope, "obvious", although the formal reason, namely Dedekind's axiom, is hidden in 312.] Here

$$f^{-1}((-\infty, \beta)) = \bigcup_{\alpha \in A} f^{-1}((-\infty, \alpha)) = \emptyset .$$

On the other hand, for $\alpha > \beta$, $f^{-1}((\alpha, \infty)) = \emptyset$, since $f^{-1}((-\infty, \alpha))$ is cofinite, as above, as $\alpha \notin A$; and so $f^{-1}((\beta, \infty)) = \bigcup_{\alpha > \beta} f^{-1}((\alpha, \infty)) = \emptyset$. It follows that f takes only the one value β. Therefore, the only continuous functions on Ω with values in $\mathbb{R}$ are constant.

Thus, in this case, the weak topology on Ω defined by the class of continuous functions on Ω is in fact the indiscrete topology — which has constants as its only continuous mappings into *any* topological space.

This rather extreme example raises the question, which we shall discuss later, whether there is some simple way of characterizing or recognizing topological spaces that allow large numbers of real-valued functions.

Amongst these "weak topologies" (they would more often be called "topologies induced by the family $\mathcal{F}$", but there are other situations in which topologies are "induced", as we shall see) is the special case of the *product topology*.

Definition 5.15. Let $(\Omega_\beta, \mathcal{G}_\beta)_{\beta \in B}$ be a family of topological spaces, indexed by the set B. The *coordinate projection* of the Cartesian product $\prod_{\beta \in B} \Omega_\beta$ on the "γth coordinate" Ω_γ is the mapping $\pi_\gamma : \prod_{\beta \in B} \Omega_\beta \longrightarrow \Omega_\gamma : f \mapsto f(\gamma)$. [In this formulation I am assuming the definition of the product that I gave in §4**A**. Thus, a point of $\prod_{\beta\in B}\Omega_\beta$ is, by definition, a function $f : B \longrightarrow \bigcup_{\beta\in B}\Omega_\beta$ such that $f(\beta) \in \Omega_\beta$ for each $\beta \in B$. I believe this is the only way you can *define* a general product; but if, intuitively, we think of it as the set of "B-tuples" $(x_\beta)_{\beta\in B}$, where of course $x_\beta = f(\beta)$ for each β, then

$$\pi_\gamma((x_\beta)_{\beta\in B}) = x_\gamma \, .]$$

The *product topology* is the topology on $\prod_{\beta\in B}\Omega_\beta$ induced by the coordinate projections.

Unravelling this definition, we see that the product topology is defined so that it has a subbase of sets of the form $\pi_\beta^{-1}(U)$, for any $\beta \in B$ and $U \in \mathcal{G}_{\Omega_\beta}$, and, therefore, a base consisting of "basic" sets of the form

$$\pi_{\beta_1}^{-1}(U_1) \cap \pi_{\beta_2}^{-1}(U_2) \cap \cdots \cap \pi_{\beta_k}^{-1}(U_k)\,, \tag{10}$$

where $\{\beta_1, \beta_2, \ldots, \beta_k\}$ may be any finite subset of the index set B and $U_i \in \mathcal{G}_{\beta_i}$ for $1 \le i \le k$. In the vaguer "B-tuple" description of the product, the basic open sets are the subsets of the form $\prod_{\beta\in B}\widehat{\Omega}_\beta$, where $\widehat{\Omega}_\beta = \Omega_\beta$ for all the indices β except for finitely many, $\beta_1, \beta_2, \ldots, \beta_k$ (where k may be any nonnegative integer), for each of which $\widehat{\Omega}_\beta$ is an open subset of $\mathcal{G}_\beta$ (in the representation (10), $\widehat{\Omega}_{\beta_i} = U_i$ for each i).

Notice that when B is finite (for instance if one has the Cartesian product $\Omega_1 \times \Omega_2$ of just two spaces) there is a base for the product topology consisting of products of open sets in the individual spaces. In $\Omega_1 \times \Omega_2$, the product topology has a base $\{U_1 \times U_2 : U_1 \in \mathcal{G}_{\Omega_1}\ \&\ U_2 \in \mathcal{G}_{\Omega_2}\}$. The naturally corresponding statement is false if B is infinite, because only *finite* intersections appear in (10).

There are a couple of significant facts about product topologies that are conveniently treated as exercises. Firstly, if $\mathcal{B}_\beta$ is a base for the topology $\mathcal{G}_{\Omega_\beta}$ for each $\beta \in B$, then sets of the form

$$\pi_{\beta_1}^{-1}(B_1) \cap \pi_{\beta_2}^{-1}(B_2) \cap \cdots \cap \pi_{\beta_k}^{-1}(B_k)\,,$$

where $B_i \in \mathcal{B}_{\beta_i}$ for each i, constitute a base for the product topology in $\prod_{\beta\in B}\Omega_\beta$. Secondly, the product topology construction is "associative" in the (more or less) obvious sense. In particular, the natural mappings between $(\Omega_1 \times \Omega_2) \times \Omega_3$, $\Omega_1 \times (\Omega_2 \times \Omega_3)$, and $\Omega_1 \times \Omega_2 \times \Omega_3$ that correspond to inserting or removing parentheses are homeomorphisms with respect to the product topologies.

Remark 5.16. The coordinate projections are continuous, because the product topology is defined to ensure they are; but, in contrast to the general weak topologies of 5.13, they are also open (compare 5.10). This is an easy exercise.

The product topology has a "universal mapping property".

Lemma 5.17. *Suppose that* Ψ *is a topological space and that* $g_\beta : \Psi \longrightarrow \Omega_\beta$ *is continuous for each* $\beta \in B$. *Then there is one and only one continuous mapping* $g : \Psi \longrightarrow \prod_{\beta\in B}\Omega_\beta$ *such that* $\pi_\beta g = g_\beta$ *for each* $\beta \in B$.

Proof. Indeed, we define, for each $x \in \Psi$, $g(x) := (g_\beta(x))_{\beta\in B}$ in the "B-tuple" notation; that is, $g(x)$ is the choice function $\phi_x : B \longrightarrow \bigcup_{\beta\in B}\Omega_\beta$ such that $\phi_x(\beta) := g_\beta(x)$ for each β. Evidently $\pi_\beta g = g_\beta$, and g is the only possible *mapping* $\Psi \longrightarrow \prod_{\beta\in B}\Omega_\beta$ with this property. g is continuous by 3.7*(b)*, since, for any subbasic open set $\pi_\beta^{-1}(U)$ (where $\beta \in B$ and $U \in \mathcal{G}_{\Omega_\beta}$), $g^{-1}(\pi_\beta^{-1}(U)) = g_\beta^{-1}(U) \in \mathcal{G}_\Psi$ from 3.3. ☐

It is common to describe this situation by a diagram such as

$$\begin{tikzcd} & \Omega_\delta \\ \Psi \arrow[ur] \arrow[r, dashed, "g"'] \arrow[dr, "g_\gamma"'] & \prod_{\beta\in B}\Omega_\beta \arrow[u, "\pi_\delta"'] \arrow[d, "\pi_\gamma"] \\ & \Omega_\gamma \end{tikzcd} \tag{11}$$

where the solid arrows indicate mappings that are *given*, the broken arrow is the mapping whose existence is being *asserted*, and the triangles are *commutative*, i.e. the same mapping results from going by either of the routes allowed by the arrows from one vertex to another; thus $\pi_\gamma g = g_\gamma$, for instance. [functoriality]

Lemma 5.18. *The product topology and the Euclidean topology on* $\mathbb{R}^2$ *coincide.* ☐

This fact can be generalized (see the exercises).

Lemma 5.19. *The mapping* $(x, y) \mapsto x - y : \mathbb{R}^2 \longrightarrow \mathbb{R}$ *is continuous.* ☐

Again this is a very small instance of a very general fact.

So far I have been considering "weak" topologies, where a topology is defined on a space by means of mappings out of it. But we can also try to define a topology by means of mappings into a space, which is a "dual" problem.

Definition 5.20. Let $\{(\Psi_\beta, \mathcal{G}_{\Psi_\beta}) : \beta \in B\}$ be a family of topological spaces, let Ω be a set, and suppose that for each $\beta \in B$ a mapping $h_\beta : \Psi_\beta \longrightarrow \Omega$ is given. The *topology induced on* Ω *by the mappings* h_β is

$$\mathcal{G}_\Omega := \{U \in \mathcal{P}(\Omega) : (\forall \beta \in B)\ \ h_\beta^{-1}(U) \in \mathcal{G}_{\Psi_\beta}\}\,. \tag{12}$$

Lemma 5.21. $\mathcal{G}_\Omega$, *as defined by* (12), *is a topology on* Ω. [Use 3.2.] ☐

Lemma 5.22. (a) *The topology* $\mathcal{G}_\Omega$ *defined by* (12) *is the finest topology such that all the mappings* h_β *are continuous.*

(b) *Let* $(\Phi, \mathcal{G}_\Phi)$ *be any topological space. Then* $g : (\Omega, \mathcal{G}_\Omega) \longrightarrow (\Phi, \mathcal{G}_\Phi)$ *is continuous if and only if* $gh_\beta : (\Psi_\beta, \mathcal{G}_{\Psi_\beta}) \longrightarrow (\Phi, \mathcal{G}_\Phi)$ *is continuous for every* $\beta \in B$.

(c) $\mathcal{G}_\Omega$ *is the only topology on* Ω *making* (b) *true for all mappings* g. ☐

Remark 5.23. Notice that, if $\mathcal{G}_\Omega$ is defined by (12), then $A \in \mathcal{P}(\Omega)$ is closed if and only if $h_\beta^{-1}(A)$ is closed in Ψ_β for each $\beta \in B$. (This follows from 3.2.)

A special case is the topology on Ω induced from the inclusion $i_\Psi : \Psi \longrightarrow \Omega$ when Ψ is a subset of Ω furnished with a topology $\mathcal{G}_\Psi$. In this case

$$\mathcal{G}_\Omega := \{U \in \mathcal{P}(\Omega) : U \cap \Psi \in \mathcal{G}_\Psi\}.$$

Thus $\mathcal{G}_\Omega$ consists of all unions of open sets of Ψ with (arbitrary) subsets of $\Omega \setminus \Psi$. For instance, if $h : \Psi \longrightarrow \Omega$ is constant (its image is a single point), then $\mathcal{G}_\Omega$ will be the discrete topology.

One reason why this topology is "unsatisfactory" is that the inclusion is not surjective; all subsets of $\Omega \setminus \Psi$ are both open and closed. This suggests considering what happens for surjections.

Remark 5.24. A surjection $h : \Psi \longrightarrow \Omega$ defines an equivalence relation $\sim$ on Ψ by the rule "$x \sim y$ means $h(x) = h(y)$". Conversely, if $\sim$ is an equivalence relation on Ψ, the *quotient of* Ψ *by* $\sim$, denoted $\Psi/\sim$, is the set of $\sim$-equivalence classes in Ψ. The *quotient map* (or *projection*) $\pi : \Psi \longrightarrow \Psi/\sim$ is the mapping defined by

$$(\forall x \in \Psi) \quad \pi(x) := [x]_\sim$$

(where $[x]_\sim$ denotes the $\sim$-equivalence class containing x). π is surjective.

In this way, any surjection out of Ψ defines an equivalence relation on Ψ, and any equivalence relation defines a surjection. We have a commutative diagram

$$\begin{tikzcd} & \Psi \arrow[dl] \arrow[dr] & \\ \Omega \arrow[rr, dashed] & & \Psi/\sim \end{tikzcd}$$

relating the map h and the quotient projection; the bottom arrow is a bijection, so that, from a set-theoretic point of view, Ω and the quotient $\Psi/\sim$ can be regarded as "essentially the same".

Definition 5.25. Let $(\Psi, \mathcal{G}_\Psi)$ be a topological space, and let $f : \Psi \longrightarrow \Omega$ be a surjection. In this case the topology $\mathcal{G}_\Omega$ defined by (12) is called the *quotient topology*. It is the finest topology such that f is continuous.

It might be natural to call the topology $\mathcal{G}_\Omega$ of (12) the "strong topology" defined by the family of mappings (h_β). This is not standard practice. In fact, I was careful to say that an *analyst* would describe the topology of 5.13 as a "weak topology", because algebraic topologists have been known (in certain special cases) to call the topology of (12) a "weak topology" instead.

Since 5.22 is a statement which just "reverses the arrows" of 5.13, it is natural to ask whether there is an object which will reverse the arrows of the universal mapping property of the product, as illustrated by (11). Such an object would be called a *coproduct* in category theory, and would be denoted by $\coprod$. And it does exist here.

Lemma 5.26. *Let* $\{(\Omega_\beta, \mathcal{G}_\beta) : \beta \in B\}$ *be a family of topological spaces. There is a topological space* Ω *(customarily called the* disjoint union *of the* Ω_β*) and there are continuous mappings* $i_\beta : \Omega_\beta \longrightarrow \Omega$ *such that, for any topological space* Φ *and any* *system* $(g_\beta)_{\beta\in B}$ *of continuous mappings* $g_\beta : \Omega_\beta \longrightarrow \Phi$, *there is a unique continuous map* $g : \Omega \longrightarrow \Phi$ *for which* $gi_\beta = g_\beta$ *for each* $\beta \in B$.

Proof. As a set, Ω is the "disjoint union" of the Ω_β. [If the sets Ω_β are already (pairwise) disjoint, i.e. no two of them have a point in common, their union is already a disjoint union. But it is entirely possible that they are not disjoint *a priori* — in the most extreme case, they might all be the same nonnull set; but we want the union of "distinct copies". The simplest way of overcoming this difficulty is to define the disjoint union by

$$\coprod_{\beta\in B} \Omega_\beta := \bigcup_{\beta\in B} (\Omega_\beta \times \{\beta\}), \tag{13}$$

because the summands on the right-hand side are certainly (pairwise) disjoint and each is in a natural one-one correspondence with Ω_β. But (partly out of laziness) I shall write Ω_β for the βth summand anyway, tacitly assuming all summands either are disjoint or have been made disjoint by some trick like the one above. Although I use the coproduct sign, people often write an ordinary union $\bigcup_{\beta\in B} \Omega_\beta$ with the verbal explanation that they mean the "disjoint union".] Then $i_\beta : \Omega_\beta \longrightarrow \Omega$ is just inclusion [or, if we accept (13), $i_\beta(x) := (x, \beta)$ for any $x \in \Omega_\beta$].

Use (12) to define a topology on Ω. In fact, in this case the open sets of Ω are precisely the (disjoint) unions $\coprod_{\beta\in B} U_\beta$, where U_β is open in Ω_β for each $\beta \in B$. Given the family (g_β) of continuous mappings as in the statement, g is necessarily defined by setting $g(x) := g_\beta(x)$ for each $\beta \in B$ and each $x \in \Omega_\beta$. It is now trivial that g is continuous. □

*****second countable

The conclusion may be expressed by the commutative diagram dual to (11)

$$\begin{tikzcd}
& \Omega_\delta \arrow[dl] \arrow[d, "i_\delta"] \\
\Psi & \coprod_{\beta\in B} \Omega_\beta \arrow[l, dashed, "g"'] \\
& \Omega_\gamma \arrow[ul, "g_\gamma"] \arrow[u, "i_\gamma"']
\end{tikzcd}$$

§6. Convergence.

The reason general topology was invented was to embrace ideas of convergence and continuity more general than just convergence of sequences in metric spaces. There are several ways of doing this.

Definition 6.1. Let $(x_n)_{n=1}^\infty$ be a sequence in the metric space (Ω, d), and suppose $x \in \Omega$. Then (x_n) *converges* (or *tends*) to x (as $n \to \infty$), denoted $x_n \to x$ or $\lim_{n\to\infty} x_n = x$ (x is the *limit* of the sequence (x_n)) if

$$(\forall \epsilon > 0)(\exists N \in \mathbb{N}) \quad n \geq N \Longrightarrow d(x_n, x) < \epsilon\,.$$

[That is to say: however close to x we should like the terms x_n to be, for instance within the distance ϵ for any positive ϵ we choose, they will be as close as that once n is sufficiently large — at least as large as some number N. Notice that N 'depends on' the choice of ϵ in the weak sense that, if we change ϵ to a smaller value, we usually need a larger value of N. For this reason people often write $N(\epsilon)$ to indicate that the N under consideration is one that works for the given ϵ.] ****uniqueness

The obvious generalization to topological spaces is the one I gave in 312:

Definition 6.2. The sequence (x_n) in the topological space $(\Omega, \mathcal{G})$ converges to $x \in \Omega$ (which may be called a *limit* of the sequence) if

$$(\forall U \in \mathfrak{N}(x))(\exists N \in \mathbb{N}) \quad n \geq N \Longrightarrow x_n \in U\,.$$

This can be interestingly reformulated. If we write $f(1/n) := x_n$ and $f(0) := x$, then the statement that $x_n \to x$ is precisely equivalent to the statement that $f : A \longrightarrow \Omega$ is continuous when $A := \{1/n; n \in \mathbb{N}\} \cup \{0\}$ is topologized as a subspace of $\mathbb{R}$.

That 6.2 is equivalent in a metric space to 6.1 is trivial. Indeed, in a metric space, the topology can be described in terms of convergent sequences, and more generally

Lemma 6.3. *Let Ω be a first countable topological space* [recall 2.13 and 2.15]. *For any $A \in \Omega$, $\mathrm{cl}(A)$ is the set of limits of sequences in A (i.e. having all their terms in A) that are convergent in Ω.*

Proof. $x \in \mathrm{cl}(A)$ if and only if every neighbourhood of x meets A (2.11*(a)*). By 2.14, there is a decreasing base (V_n) of neighbourhoods of x. Thus $x \in \mathrm{cl}(A)$ if and only if $V_n \cap A \notin \emptyset$ for each n. If that is so, then, for each n, there is some $x_n \in V_n \cap A$, and the sequence (x_n) (in A) converges to x. (Indeed, if $M \in \mathfrak{N}(x)$, there is some $N \in \mathbb{N}$ such that $V_N \subseteq M$, and then

$$n \geq N \Longrightarrow x_n \in V_n \subseteq V_N \subseteq M\,,$$

satisfying 6.2. On the other hand, if (x_n) is a sequence in A such that $x_n \to x$, then, for any $U \in \mathfrak{N}(x)$, there is (from 6.2) some N such that $x_N \in U$, and so $U \cap A \notin \emptyset$. So $x \in \mathrm{cl}(A)$. ☐

It is easily seen that, if $f : \Omega \longrightarrow \Psi$ is continuous at $x \in \Omega$ and (x_n) is a sequence in Ω such that $x_n \to x$, necessarily $f(x_n) \to f(x)$ in Ψ.

Definition 6.4. Let Ω and Ψ be topological spaces, $f : \Omega \longrightarrow \Psi$, and $x \in \Omega$. f is *sequentially continuous* at x if, for any sequence (x_n) in Ω with $x_n \to x$, $f(x_n) \to f(x)$ in Ψ.

Lemma 6.5. *If f is continuous at x, then f is sequentially continuous at x.* ☐

The point is that, in some sense and in some circumstances, sequential continuity is easier to check than continuity.

Lemma 6.6. *Suppose that* $x \in \Omega$ *has a countable base of neighbourhoods, and that* $f : \Omega \longrightarrow \Psi$ *is sequentially continuous at* x*. Then* f *is continuous at* x*.*

Proof. Let C be a closed subset of Ψ with $f(x) \notin C$, $x \notin f^{-1}(C)$. If $x \in \mathrm{cl}_\Omega(f^{-1}(C))$, then, by 6.3, there is a sequence (x_n) in $f^{-1}(C)$ such that $x_n \to x$. By hypothesis, $f(x_n) \to f(x)$. But this is absurd, as $\Psi \setminus C$ is an open set containing $f(x)$ but none of the terms $f(x_n)$. Hence $x \notin \mathrm{cl}_\Omega(f^{-1}(C))$, and this proves continuity at x, by 3.6. □

However, the condition that there should be a countable base of neighbourhoods really cannot be dispensed with. Here is a rather over-subtle example.

Example 6.7. Let Ω be the set of all continuous functions $[0,1] \longrightarrow [0,1]$. Thus Ω is (by definition!) a subset of the product space $\prod_{t\in[0,1]} [0,1]$, which may be given the product topology. Give Ω the subspace topology of the product topology.

Define the mapping $\phi : \Omega \longrightarrow [0,1]$ by $\phi(f) := \int_0^1 f(t)\, dt$. The integral will make sense, because f is a continuous real-valued function on $[0,1]$.

Suppose that $g_n \to g$ in Ω. This implies that $g_n(t) \to g(t)$ for each $t \in [0,1]$ (the mapping $f \mapsto f(t)$ is the coordinate projection on the tth coordinate, cf. 5.15, and so is continuous); that is, $g_n \to g$ "pointwise". Now it is an important theorem from integration theory (I shan't prove it here, but it is the *dominated convergence theorem*, whose essential hypothesis is that the functions g_n all take values in $[0,1]$) that in these circumstances $\int g_n \to \int g$. Thus the mapping $\phi : \Omega \longrightarrow \mathbb{R}$ is sequentially continuous.

But it is not continuous. Indeed, if $g \in \Omega$, there is a base of neighbourhoods of g consisting of sets of the form $\{f \in \Omega : f(t_i) \in U_i\,,\ 1 \le i \le k\}$ for all the various choices of $k \in \mathbb{N}$, of finite subsets $\{t_1, t_2, \ldots, t_k\}$ of $[0,1]$, and of neighbourhoods U_i of $g(t_i)$ for each i. However, you can easily see that there will be functions $f \in \Omega$ such that $f(t_i) = g(t_i)$ for $1 \le i \le k$ but $\int f$ is as close as we like to either 1 or 0. Thus the topology on Ω that we have prescribed (the "topology of pointwise convergence") is just too weak for ϕ to be continuous.

Although this may seem a rather unnecessarily highbrow example, it is interesting because it puts on view a situation where sequential continuity arises *naturally*. In fact it is related to some important results in functional analysis.

Notice that in the topology of Ω no point can have a countable base of neighbourhoods (either because of 6.6, or by direct argument).

In rather vague terms, we might say that the idea of convergence of sequences is enough to express all the concepts of topology provided that we restrict attention to first-countable spaces. The question arises whether we can modify the notion of convergence so as to remove the need for first-countability. This was originally done by E. H. Moore and H. L. Smith as early as 1922, although, both in its details and as a tool in topology, it was improved gradually over quite a long time. But there is an alternative approach due to Henri Cartan in 1937, and it has the advantage of giving snappy proofs of some otherwise nasty theorems. So I shall present both.

The Moore-Smith approach begins by generalizing the idea of a sequence.

Definition 6.8. Let $(D, \le)$ be a partially ordered set. It is a *directed set*, and the partial order $\le$ is called a *direction* (on D), if

$$(\forall a, b \in D)(\exists c \in D) \quad a \leq c \;\; \& \;\; b \leq c\,.$$

We commonly just say "the directed set D", understanding that the direction will be denoted by $\leq$ unless there is some reason to choose another notation. Incidentally, it is often convenient to write $a \geq b$ to mean the same as $b \leq a$.

A subset E of the directed set $(D, \leq)$ is *cofinal* if, for all $m \in D$, there is an element $n \in E$ such that $m \leq n$. [Notice the word is "cofinal", not "cofinite"! The idea is that a cofinal subset contains "arbitrarily large" elements of D.]

Let Ω be any set. A *net* (or *generalized sequence*) in Ω is a function $x : D \longrightarrow \Omega$ together with a direction on D. (One may say simply that it is a function from a directed set to Ω.) As with sequences, one customarily writes x_a instead of $x(a)$, and $(x_a)_{a \in D}$ instead of $x : D \longrightarrow \Omega$.

Clearly $\mathbb{N}$ is a directed set and a sequence in the usual sense is a net when $\mathbb{N}$ is given the standard direction. This suggests the definition:

Definition 6.9. Let Ω be a topological space and $(x_a)_{a \in D}$ a net in Ω. If $x \in \Omega$, we say that x_a or (x_a) *tends to* x (for $a \in D$) or *converges to* x, or that x is the *limit* of the net, $\lim_{a \in D} x_a = x$, or that $x_a \underset{D}{\to} x$ (or various other notations), if

$$(\forall M \in \mathfrak{N}(x))(\exists m \in D) \quad m \leq d \in D \Longrightarrow x_d \in M\,.$$

This is of course just the usual definition for sequences, reformulated for nets. [It is worth noting that it can only apply if $D \neq \emptyset$. I did not need explicitly to exclude the "empty net" in the definitions.]

Lemma 6.10. *Let A be a subset of the topological space Ω. Then a point $x \in \Omega$ belongs to $\mathrm{cl}(A)$ if and only if there is a net in A which converges (in Ω) to x.*

Proof. If $(x_a)_{a \in D}$ is a net in A convergent to x, let $M \in \mathfrak{N}(x)$. There is some $m \in D$ such that $m \leq d \in D \Longrightarrow x_d \in M$. Thus $M \cap A \neq \emptyset$. As this holds for any neighbourhood of x, $x \in \mathrm{cl}(A)$.

Conversely, suppose $x \in \mathrm{cl}(A)$. The crucial observation is that $\mathfrak{N}(x)$ itself is a directed set, if we define $M_1 \leq M_2$ to mean $M_1 \supseteq M_2$. This is clearly a partial order; it is a direction because, for M_3 and M_4 in $\mathfrak{N}(x)$, $M_3 \leq M_3 \cap M_4 \in \mathfrak{N}(x)$ and $M_4 \leq M_3 \cap M_4$. For each $M \in \mathfrak{N}(x)$, choose $x_M \in M \cap A \neq \emptyset$. Then $(x_M)_{M \in \mathfrak{N}(x)}$ is a net in A which converges to x. □

Notice that the Axiom of Choice has been used in this proof.

Lemma 6.11. *A mapping $f : \Omega \longrightarrow \Psi$ between topological spaces is continuous at $x \in \Omega$ if and only if, for any net $(x_a)_{a \in D}$ in Ω that converges to x, the net $(f(x_a))_{a \in D}$ in Ψ converges to $f(x)$.*

Proof. That the condition is *necessary* is an easy exercise.

Suppose, therefore, that it is satisfied. Let C be a closed set in Ψ such that $f(x) \notin C$. Now, if $x \in \mathrm{cl}_\Omega(f^{-1}(C))$, then by 6.10 there is a net $(x_a)_{a \in D}$ in $f^{-1}(C)$ that converges to x. But then, by hypothesis, $f(x_a) \to f(x)$, and each

$f(x_a)$ is in C. By 6.10, $f(x) \in C$, contrary to assumption. Consequently, $x \notin \mathrm{cl}_\Omega(f^{-1}(C))$. But this means that f is continuous at x, by 3.6. $\square$

For some purposes it is desirable to have an idea of "subnet" generalizing the concept of "subsequence". To obtain a definition that is sufficient for the aims in view is not quite as trivial a matter as you might expect.

Definition 6.12. Let $(x_a)_{a\in D}$ be a net (in a set Ω). A *subnet* of (x_a) is a function $\phi : E \longrightarrow D$, where E is a directed set, such that

$$(\forall d \in D)(\exists e \in E) \quad (c \in E \ \& \ e \leq c) \Longrightarrow d \leq \phi(c) . \tag{14}$$

The subnet would usually be denoted $(x_{\phi(b)})_{b\in E}$.

The condition (14) ensures "cofinality is preserved". It is true, in particular, if E is a cofinal subset of D and ϕ is just the inclusion. At first sight this might seem a more natural definition of a subnet, but there are technical reasons (not worth discussing here) why one sometimes needs the more general definition.

Definition 6.13. Let $(x_d)_{d\in D}$ be a net in the topological space Ω. A *cluster point* of this net is a point $x \in \Omega$ such that, for every $M \in \mathfrak{N}(x)$, $\{d \in D : x_d \in M\}$ is cofinal in D; i.e., for any $a \in D$, there is some $d \in D$ with $a \leq d$ and $x_d \in M$.

It should be emphasized that this is a concept that is tied to *nets* (and, of course, to sequences as a special case of nets).

Let me now proceed to Cartan's method of discussing convergence.

Definition 6.14. Let Ω be a nonnull set and $\mathfrak{Q}$ a class of subsets of Ω such that $\Omega \in \mathfrak{Q}$. [By far the most important case is when $\mathfrak{Q} = \mathcal{P}(\Omega)$.] A *filter* in $\mathfrak{Q}$ is a nonempty subclass $\mathfrak{F}$ of $\mathfrak{Q}$ such that

(a) $\emptyset \notin \mathfrak{F}$,

(b) if $A \in \mathfrak{Q}$ and $A \supseteq B \in \mathfrak{F}$, then $A \in \mathfrak{F}$,

(c) if $A, B \in \mathfrak{F}$, then $A \cap B \in \mathfrak{F}$.

If $\mathfrak{Q} = \mathcal{P}(\Omega)$, we speak of a filter in Ω. For most classes $\mathfrak{Q}$, *(b)* and *(c)* together may be difficult to satisfy unless the intersection of two members of $\mathfrak{Q}$ always belongs to $\mathfrak{Q}$. The only example that will really interest us is where $\mathfrak{Q} = \mathcal{F}$, the class of all closed sets.

If $\mathfrak{F}_1, \mathfrak{F}_2$ are filters in $\mathfrak{Q}$ (or in Ω), we say that $\mathfrak{F}_2$ *refines* $\mathfrak{F}_1$, or that $\mathfrak{F}_2$ is a *refinement* of $\mathfrak{F}_1$, if $\mathfrak{F}_1 \subseteq \mathfrak{F}_2$, i.e. every set belonging to $\mathfrak{F}_1$ also belongs to $\mathfrak{F}_2$.

You may have met filters in other contexts where the condition *(a)* is not imposed. It has the important consequence, with *(c)*, that *any finite intersection of members of* $\mathfrak{F}$ *is nonempty*. Condition *(b)* implies (as $\mathfrak{F}$ is nonempty) that $\Omega \in \mathfrak{F}$; it also means that, if $\mathfrak{F}_2$ strictly refines $\mathfrak{F}_1$ (i.e. refines it and is not the same), then, for any $A \in \mathfrak{F}_1$, there must be a $B \in \mathfrak{F}_2$ such that $B \subset A$, where $\subset$ denotes strict inclusion. There is a coarsest possible filter, consisting of Ω alone.

Because of *(b)*, a filter may often be specified by a small selection of its members:

Definition 6.15. A *filter base* in $\mathfrak{Q}$ is a subclass $\mathfrak{B}$ of $\mathfrak{Q}$ such that

(a) $\emptyset \notin \mathfrak{B} \neq \emptyset$,

(b) if $A, B \in \mathfrak{B}$, there exists $C \in \mathfrak{B}$ such that $C \subseteq A \cap B$.

If $\mathfrak{B}$ is a filter base in $\mathfrak{Q}$, the *filter in $\mathfrak{Q}$ generated by* $\mathfrak{B}$ is

$$\mathfrak{F}(\mathfrak{B}) := \{E \in \mathfrak{Q} : (\exists B \in \mathfrak{B})\ B \subseteq E\},$$

which is readily seen to be a filter. $\mathfrak{B}$ is described as a *base for the filter* $\mathfrak{F}(\mathfrak{B})$ *in* $\mathfrak{Q}$. If $\mathfrak{B}_1, \mathfrak{B}_2$ are filter bases in $\mathfrak{Q}$, we say that $\mathfrak{B}_2$ *refines* $\mathfrak{B}_1$ if $\mathfrak{F}(\mathfrak{B}_2)$ refines $\mathfrak{F}(\mathfrak{B}_1)$.

It is easily seen that $\mathfrak{F}(\mathfrak{B}_1) \subseteq \mathfrak{F}(\mathfrak{B}_2)$ if and only if, for every $B_1 \in \mathfrak{B}_1$, there is some $B_2 \in \mathfrak{B}_2$ such that $B_2 \subseteq B_1$. The word "refinement" is quite appropriate.

Lemma 6.16. *If Ω is a topological space and $x \in \Omega$, $\mathfrak{N}(x)$ is a filter in Ω. A base of neighbourhoods at x is exactly the same as a base for the filter $\mathfrak{N}(x)$.*

Proof. See 2.8*(a), (b), (c).* □

In fact, $\mathfrak{N}(x)$ is usually called the *neighbourhood filter at x*. If $\mathfrak{Q}$ is taken to be the class of closed subsets of Ω, then the *closed* neighbourhoods of x form a filter in $\mathfrak{Q}$. (However, so far there need not be many closed neighbourhoods.)

Definition 6.17. Let Ω be a topological space, $x \in \Omega$, and $\mathfrak{F}$ a filter in Ω. We say that $\mathfrak{F}$ *converges* or *tends* to x, or that x is the *limit* of $\mathfrak{F}$, and we write $\mathfrak{F} \to x$ and so on, if $\mathfrak{F}$ refines $\mathfrak{N}(x)$, i.e. $\mathfrak{N}(x) \subseteq \mathfrak{F}$. If $\mathfrak{B}$ is a filter base in Ω, we say that $\mathfrak{B}$ converges to x (with the same synonyms) if $\mathfrak{F}(\mathfrak{B})$ converges to x.

This is obviously a more "abstract" definition than for nets, and at first sight it is difficult to see how it relates to the familiar idea of convergence of sequences. Nets are (in retrospect) a natural generalization of sequences, and often proofs in terms of nets, valid in spaces that are not first countable, can be constructed by a simple rephrasing of arguments with sequences, whereas filters are harder to grasp.

Lemma 6.18. *Let $(x_d)_{d\in D}$ be a net in a set Ω, where D is of course a directed set. Then the "segments" $\Lambda(c) := \{x_d : c \le d \in D\}$ form a filter base $\mathfrak{B}$ in Ω as c varies over D. Furthermore, if Ω is a topological space and $x \in \Omega$, $x_d \underset{D}{\to} x$ if and only if $\mathfrak{B} \to x$.*

Proof. Since $x_c \in \Lambda(c)$, 6.15*(a)* is immediate. If $c_1, c_2 \in D$, then there exists $c_3 \in D$ such that $c_1 \le c_3 \ge c_2$, since D is directed. If $y \in \Lambda(c_3)$, then $y = x_d$ for some $d \ge c_3$, and then $d \ge c_2$ and $d \ge c_1$ also, so that $y = x_d \in \Lambda(c_1) \cap \Lambda(c_2)$. Thus $\Lambda(c_3) \subseteq \Lambda(c_1) \cap \Lambda(c_2)$, which establishes 6.15*(b).*

To say $\mathfrak{B} \to x$ is to assert that, for any $M \in \mathfrak{N}(x)$, $M \in \mathfrak{F}(\mathfrak{B})$; in turn that says there exists some $B \in \mathfrak{B}$ with $B \subseteq M$, i.e. that, for some $c \in D$, $\Lambda(c) \subseteq M$. But "$(\forall M \in \mathfrak{N}(x))(\exists c \in D)\ \Lambda(c) \subseteq M$" is just 6.13. □

Thus a net gives rise in a natural way to a filter base. On the other hand, a filter base $\mathfrak{B}$ (or, of course, a filter) is a directed set with respect to "reverse inclusion", $A \le B$ meaning $A \supseteq B$. This follows from 6.15*(b)* (or 6.14*(c)*). Since each member of $\mathfrak{B}$ is nonempty, we may choose an element $x_B \in B$ for each $B \in \mathfrak{B}$, and the result will be a net $(x_B)_{B\in\mathfrak{B}}$. Such a net is in effect a choice function for $\mathfrak{B}$; let me call it a net "associated to $\mathfrak{B}$".

Lemma 6.19. *The filter base $\mathfrak{B}$ in the topological space Ω converges to the point x if and only if every net associated to $\mathfrak{B}$ converges to x.*

Proof. Suppose $\mathfrak{B} \to x$. Let $M \in \mathfrak{N}(x)$. There is $B_0 \in \mathfrak{B}$ such that $B_0 \subseteq M$; thus, for *any* associated net $(x_B)_{B \in \mathfrak{B}}$, $x_B \in B \subseteq B_0 \subseteq M$ when $B \geq B_0$. The definition 6.9 is satisfied.

Conversely, suppose that $\mathfrak{B} \not\to x$. Then there is some $M \in \mathfrak{N}(x)$ such that, for every $B \in \mathfrak{B}$, $B \not\subseteq M$, that is, $B \setminus M \neq \emptyset$. But $\{B \setminus M : B \in \mathfrak{B}\}$ is, therefore, a filter base (check this), and refines $\mathfrak{B}$. Any net associated to it will also be associated to $\mathfrak{B}$, but will not converge to x. □

You see that in dealing with nets we are repeatedly forced to use the axiom of choice (for instance, the proof just given has to assume that there *is* an associated net for the filter base $\{B \setminus M : B \in \mathfrak{B}\}$). Crudely, one might say that filter bases are inchoate nets before the actual values of the terms of the net are fixed.

Lemma 6.20. *Let $\mathfrak{B}$ be a filter base in a set Ω, and $f : \Omega \longrightarrow \Psi$. Then*

$$f(\mathfrak{B}) := \{f(B) : B \in \mathfrak{B}\}$$

is a filter base in Ψ. □

Lemma 6.21. *Let $f : \Omega \longrightarrow \Psi$ be a mapping between topological spaces. Then f is continuous at $x \in \Omega$ if and only if, for every filter base $\mathfrak{B}$ in Ω such that $\mathfrak{B} \to x$ in Ω, the image filter base $f(\mathfrak{B})$ converges in Ψ to $f(x)$.* □

§7. Separation axioms.

Many early results in topology involved conditions on abstract topological spaces sufficient to ensure that familiar facts from $\mathbb{R}^n$ would generalize. For instance, the cofinite topology on an infinite set Ω has no non-constant real-valued continuous functions (see 5.14), whereas one feels that a "respectable" space should have plenty. Can one invent reasonable conditions that imply there will be many continuous real-valued functions? As usual, the motivation was the desire to generalize results to more abstract and general situations.

Definition 7.1. A topological space Ω is a *T_0-space* or *is T_0* if, whenever $x, y \in \Omega$ and $x \neq y$, there is an open set U in Ω such that $U \cap \{x, y\}$ is a singleton; in other words, U contains one of x and y but not the other, and in principle we cannot choose which is which. [Kolmogorov]****Eng p. 70

Ω is a *T_1-space* or *is T_1* if, whenever $x, y \in \Omega$ and $x \neq y$, there is an open set U such that $x \in U$ and $y \notin U$. (Here we can choose freely which point will be in U and which outside it.) [Riesz, 1907]

The indiscrete topology is the only one we have met that is not T_0. 1.5 gives an example of a space that is T_0 but not T_1 (there is no open set except Ω itself that contains the "distinguished" point a.) Evidently T_1 implies T_0.

Lemma 7.2. *A space Ω is T_1 if and only if all singletons are closed.* [This is usually expressed inaccurately as "points are closed".] $\square$

Definition 7.3. A space Ω is T_2 if, whenever $x, y \in \Omega$ and $x \neq y$, there are open sets U, V in Ω such that $x \in U$, $y \in V$, and $U \cap V = \emptyset$. [Hausdorff, 1914]

This is overwhelmingly the most commonly cited separation axiom, and such an Ω would most often be called a *Hausdorff space*. ***Eng p.58

Definition 7.4. A space Ω is T_3 if

(a) it is T_1 (so that points are closed) and

(b) for any point x and closed set F not containing x, there are open sets U, V such that $x \in U$, $F \subseteq V$, and $U \cap V = \emptyset$.

It is clear that a T_3 space is Hausdorff, $\mathrm{T}_3 \Longrightarrow \mathrm{T}_2$. It is often alternatively described as *regular*. (I need scarcely stress that this word is over-used.) Kelley, in his book, draws a distinction, making "regular" mean *(b)* and "T_3" mean *(a)* & *(b)*. Although this is not silly (cf. the next lemma, where the T_1 condition is an irrelevancy), I doubt whether it is worthwhile, and it has not caught on elsewhere.

Lemma 7.5. Ω *is regular if and only if points are closed and every point has a base of neighbourhoods consisting of closed sets* [a "base of closed neighbourhoods"].

Proof. Let $M \in \mathfrak{N}(x)$. By definition, there is an open set O such that $x \in O \subseteq M$. Let $F := \Omega \setminus O$, which is closed; $x \notin F$. Hence, if Ω is regular, there are open sets U and V such that $x \in U$, $F \subseteq V$, and $U \cap V = \emptyset$, or in other terms

$$x \in U \subseteq \Omega \setminus V \subseteq \Omega \setminus F = O\,.$$

Hence, $\Omega \setminus V$ is a closed neighbourhood of x and is included in O. Any neighbourhood of x includes a closed neighbourhood; so closed neighbourhoods form a base of neighbourhoods. The argument reverses rather trivially. $\square$

Lemma 7.6. *Any subspace of a T_0-space is T_0. Any subspace of a T_1-space is T_1. Any subspace of a T_2-space is T_2. Any subspace of a T_3-space is T_3.* $\square$

Statements of this kind are usually made in the form "Property T_i is *hereditary*", i.e. is "inherited" by all subspaces.

Definition 7.7. A space Ω is T_4 if

(a) it is T_1, and

(b) for any pair E, F of disjoint closed sets in Ω, there are open sets U, V such that $E \subseteq U$, $F \subseteq V$, and $U \cap V = \emptyset$.

It is clear that *(b)* need only hold for *nonnull* disjoint closed sets E and F to be true in absolute generality; that is, if E or F is empty, it is true anyway.

Rather as before, $\mathrm{T}_4 \Longrightarrow \mathrm{T}_3$; T_4 spaces are alternatively described as *normal*; and again, Kelley calls a space "normal" if it just satisfies *(b)*. But normality is *not* in principle hereditary, the problem being that, if two subsets of the subspace Ψ of Ω are closed and disjoint, they may not extend to closed subsets of Ω that are also disjoint. It is not easy to find examples of this possibility, but they do exist. On the other hand,

Lemma 7.8. *If Ψ is a closed subspace of a normal space Ω, then Ψ is normal.* □

People sometimes say that *normality is hereditary for closed subspaces*.

The Ts do continue further, but "T_5" and "T_6" are rarely used in preference to the verbal names.

Definition 7.9. A space Ω is T_5 if it is *hereditarily normal*, i.e. if every subspace of Ω is normal.

Despite appearances, this is indeed a "separation axiom", as *(b)* below shows.

Lemma 7.10. *The following conditions are equivalent.*

(a) *Ω is hereditarily normal.*

(b) *For any two subsets A and B such that*

$$A \cap \operatorname{cl}(B) = \emptyset = B \cap \operatorname{cl}(A)\,, \tag{15}$$

there exist open sets U, V such that $A \subseteq U$, $B \subseteq V$, and $U \cap V = \emptyset$.

(c) *Every open subspace of Ω is normal (in the subspace topology).*

(15) is sometimes expressed by saying that A and B are *separated* in Ω. It means, of course, that neither contains any accumulation points of the other.

Proof. Certainly *(a) $\Longrightarrow$ (c)*. Now suppose that Ω satisfies *(c)*, and let A and B be separated subsets of Ω. Then $\Psi := \Omega \setminus (\operatorname{cl}(A) \cap \operatorname{cl}(B))$ is open in Ω, so is normal in the subspace topology; but, from (15), $B \subseteq \Omega \setminus \operatorname{cl}(A)$, $A \subseteq \Omega \setminus \operatorname{cl}(B)$, and, by 5.5,

$$\operatorname{cl}_\Psi(A) \cap \operatorname{cl}_\Psi(B) = \operatorname{cl}_\Omega(A) \cap \Psi \cap \operatorname{cl}_\Omega(B) \cap \Psi = \emptyset\,;$$

hence, there exist *relatively* open sets $U, V \in \mathcal{G}_\Psi$ such that $U \supseteq \operatorname{cl}_\Psi(A) \supseteq A$ and $U \supseteq \operatorname{cl}_\Psi(B) \supseteq B$ and $U \cap V = \emptyset$. However, by 5.3, U and V are open in Ω. This shows *(c) $\Longrightarrow$ (b)*. [In fact, $\operatorname{cl}_\Psi(A) = \operatorname{cl}_\Omega(A) \setminus \operatorname{cl}_\Omega(B)$ and likewise for $\operatorname{cl}_\Psi(B)$.]

Suppose now that *(b)* holds, that Ψ is a subspace of Ω, and that A, B are *relatively* closed disjoint subsets of Ψ. Then, by 5.5,

$$\begin{aligned} A \cap \operatorname{cl}_\Omega(B) &= (A \cap \Psi) \cap \operatorname{cl}_\Omega(B) = A \cap (\Psi \cap \operatorname{cl}_\Omega(B)) \\ &= A \cap \operatorname{cl}_\Psi(B) = A \cap B = \emptyset\,, \end{aligned}$$

and symmetrically $B \cap \operatorname{cl}_\Omega(A) = \emptyset$. Thus, from *(b)*, there are open sets $U, V \in \mathcal{G}_\Omega$ such that $A \subseteq U$, $B \subseteq V$, and $A \cap B = \emptyset$. But now $A \subseteq U \cap \Psi \in \mathcal{G}_\Psi$, $B \subseteq V \cap \Psi \in \mathcal{G}_\Psi$, and $(U \cap \Psi) \cap (V \cap \Psi) = \emptyset$. So Ψ is normal. □

There is also a condition T_6, or *perfectly normal*, which is a little more complicated and will be better motivated later. The first question that naturally arises is whether these conditions are often satisfied. They were in fact observed as successively more demanding properties enjoyed by metric spaces.

Proposition 7.11. *If Ω is a metric space, it is T_5.* [It is even T_6.]

It will help in the proof if I make a short digression here (which will also be relevant later).

Remark 7.12. Let A be a subset of $\mathbb{R}$. A number $l \in \mathbb{R}$ is a *lower bound* for A if

$$(\forall a \in A) \quad l \leq a\,.$$

If A has a lower bound (say l), we say it is *bounded below* (by l).

It is important to appreciate that l is to be an honest real number here (not a "symbolic" or "ideal" construct like "$-\infty$"). For instance, $\mathbb{R}$ itself is not bounded below as a subset of $\mathbb{R}$; nor is $\mathbb{Z}$, nor is $\mathbb{Q}$; but $\mathbb{N}$ is bounded below by 1, or by 0, or by $1/\pi$, or by $-\sqrt{2}$, and the "open interval" $(0,1)$ is bounded below by 0 or by any negative number. The null set is bounded below by any real number whatsoever.

We shall need the property of the real numbers usually called *Dedekind's axiom*:

if A is a *nonempty* subset of $\mathbb{R}$ and is bounded below, it has a GREATEST lower bound.

To clarify this: if A is bounded below, its lower bounds form a nonnull subset of $\mathbb{R}$. This subset *must* contain a largest element. (Thus, for instance, it cannot consist of all negative numbers, because there is no greatest negative number; if $\alpha < 0$, then $\alpha < \frac{1}{2}\alpha < 0$, so α is not a candidate for the greatest negative number.)

Dedekind's axiom is in a sense the characteristic property of real numbers; we shall see later, if it is not obvious anyway, that it says they have no "gaps". If you think of real numbers as "infinite decimals" you can easily convince yourself in intuitive terms that it is true; but a formal proof in those terms is difficult — indeed that definition of the reals makes almost everything messy to prove.

If A is a nonempty subset of $\mathbb{R}$ and is bounded below, we denote its greatest lower bound by $\inf A$ (the infimum of A). An older and less pretentious notation is $\text{g.l.b.}\,A$. If A has a least element, then $\inf A$ is that least element; but if, as for $(0,1)$, there is no least element, then $\inf A$ is the best we can do — it is less than each element of A, and is the largest number with that property. For instance

$$\inf(0,1) = 0\,, \quad \inf\{x \in \mathbb{R} : x > 0 \ \& \ x^3 > 7\} = \sqrt[3]{7}\,.$$

Definition 7.13. Let (Ω, d) be a metric space; suppose that $x \in \Omega$, $\emptyset \neq A \subseteq \Omega$. The set $\{d(x,a) : a \in A\}$ is nonnull and bounded below by 0, so has an infimum. We call this infimum $d(x,A) := \inf\{d(x,a) : a \in A\}$. (This notation is quite standard; the obvious ambiguity can never cause confusion.) Then:

Lemma 7.14. (a) $d(x,A) = 0$ *if and only if* $x \in \text{cl}_\Omega(A)$.

(b) *If* $x, y \in \Omega$, *then* $|d(x,A) - d(y,A)| \leq d(x,y)$. *In particular,* $d(x,A)$ *is continuous as a function of* x.

Proof. *(a)* $\inf\{d(x,a) : a \in A\} = 0$ if and only if no positive number is a lower bound for $\{d(x,a) : a \in A\}$; for any $\epsilon > 0$, there is some $a \in A$ ("depending on a", i.e. it may be necessary to choose a different a if the value of ϵ is changed) such that $d(x,a) < \epsilon$. But this says that $B(x;\epsilon) \cap A \neq \emptyset$ for any $\epsilon > 0$, which is equivalent to $x \in \text{cl}_\Omega(A)$.

(b) Suppose $\alpha > d(x,A)$. Then there is some $a \in A$ with $d(x,a) < \alpha$, and so $d(y,a) \leq d(x,y) + d(x,a) < d(x,y) + \alpha$; hence $d(y,A) < d(x,y) + \alpha$, and

$d(y,A) - d(x,y) < \alpha$. This is true for any $\alpha > d(x,A)$. Consequently*,

$$d(y,A) - d(x,y) \le d(x,A)\,, \quad d(y,A) - d(x,A) \le d(x,y)\,.$$

This, and the symmetrical inequality with x and y interchanged, prove *(b)*. The deduction that $d(x,A)$ is continuous in x is trivial. $\square$

[*This is a common form of reasoning, and I have given it more or less as it would usually appear. But the full argument is this. If, in fact, $d(y,A) - d(x,y) > d(x,A)$, we can take $\alpha := d(y,A) - d(x,y)$, and the proof shows that

$$d(y,A) - d(x,y) < \alpha = d(y,A) - d(x,y)\,.$$

This is absurd, and the conclusion must be that $d(y,A) - d(x,y) \le d(x,A)$.]

Lemma 7.15. *Suppose that Ω is a topological space and that $f, g : \Omega \longrightarrow \mathbb{R}$ are continuous. Then $f - g : \Omega \longrightarrow \mathbb{R} : x \mapsto f(x) - g(x)$ is also continuous.*

Proof. Define $h : \Omega \longrightarrow \mathbb{R}^2 : x \mapsto (f(x), g(x))$. This is continuous by 5.17. Then the mapping $(\alpha, \beta) \mapsto \alpha - \beta : \mathbb{R}^2 \longrightarrow \mathbb{R}$ is continuous by 5.19; so the composition of these two maps is continuous by 3.9*(a)*. $\square$

Proofs of this kind are very common. However, it would not be difficult to prove the statement directly.

We can now return to the proof of 7.11. It will be clear that, without our digression, the argument would be quite messy, although possible.

Proof of 7.11. Any subspace of a metric space is also a metric space (with a metric that defines the subspace topology; see 5.7). So it suffices to show that *a metric space is normal*. Let A and B be disjoint nonnull closed sets in the metric space Ω. Define $f : \Omega \longrightarrow \mathbb{R} : x \longrightarrow d(x,A) - d(x,B)$. This is a continuous function (by 7.14*(b)* and 7.15). Thus $U := \{x \in \Omega : f(x) < 0\}$ and $V := \{x \in \Omega : f(x) > 0\}$ are disjoint open sets in Ω. But, by 7.14*(a)*, $A \subseteq U$ and $B \subseteq V$. $\square$

If A or B is empty, 7.11 remains true, trivially so.

The principal reason why normality is an interesting condition is *Urysohn's lemma*. This is the first place where we create a continuous mapping that is not just an automatic concomitant of some abstract construction.

Theorem 7.16. (Urysohn's lemma.) *Let E and F be disjoint nonempty closed sets in a normal space Ω. Then there is a continuous function $f : \Omega \longrightarrow [0,1]$ such that*

$$f(E) = \{0\} \quad \textit{and} \quad f(F) = \{1\}\,.$$

[Incidentally, $[0,1]$ with its topology as a subspace of $\mathbb{R}$ is customarily denoted by I. I cannot understand why the authors of recent elementary textbooks believe their readers cannot cope with calling the set of integers $\mathbb{Z}$. Although a name is just a name, it seems pointless to change two that have been accepted for several generations.]

Proof. *(a)* Normality of Ω is equivalent to the property that, for any closed nonempty set E and any open $W \supseteq E$, there is an open U such that $A \subseteq U$ and

$\mathrm{cl}(U) \subseteq W$. (In 7.7, take $F := \Omega \setminus W$. Then $\mathrm{cl}(U) \subseteq \Omega \setminus V \subseteq \Omega \setminus F = W$, since $\Omega \setminus V$ is closed.) Choose one such set and call it $U(A, W)$.

(b) Let D be the set of *dyadic rational* numbers in $[0,1]$, that is, of numbers whose binary expansions, consisting of 0s and 1s, terminate. Our aim is to construct, for each $\xi \in D$, an open set U_ξ. We do this by induction on the length of the binary expansion of ξ, which I shall call the length of ξ; it is the (least) number of places (after the point) beyond which only 0s occur. We also wish to arrange that

$$\text{if } \xi < \eta \text{ in } D, \text{ then } \mathrm{cl}(U_\xi) \subseteq U_\eta . \tag{16}$$

The only numbers in $[0,1]$ of length zero are 0 and 1. Let $U_0 := U(E, \Omega \setminus F)$ and $U_1 := \Omega \setminus F$.

Now suppose U_α has been defined for all $\alpha \in D$ of length *less* than $k \in \mathbb{N}$, and that (16) is satisfied for those dyadic rationals at least. Let $\beta \in D$ be of length k. The last nonzero digit (in the kth place) in the expansion of β must be 1 (if it were 0 it would not count towards the length!) and so there are two possibilities: the last two digits in the expansion of β may be 01 or they may be 11.

Let $\alpha \in D$ be the immediate predecessor of β in the class of members of D of length not exceeding k; α must be of length less than k (it is obtained by deleting the final "1" from β, and, if $k > 1$, will have length exactly $k-1$ only if the last two digits of β are 11). Let α' be the immediate successor of α in the class of members of D of length *less* than k (for instance, if $k = 4$, the immediate successor of length less than 4 of $0 \cdot 011$ is $0 \cdot 100$, and $0 \cdot 01$ has immediate successor $0 \cdot 011$). Then $\alpha' > \beta$, and indeed α' is the immediate successor of β in the class of dyadic numbers of length not exceeding k. U_α and $U_{\alpha'}$ have already been defined. Let

$$U_\beta := U(\mathrm{cl}(U_\alpha), U_{\alpha'}),$$

in the notation of *(a)*. Then, by *(a)*, (16) holds for α, β and for β, a'. In this way U_β may be defined for all β of length k, and (16) remains true for ξ and η of length *not exceeding* k, since it suffices that it should hold for immediate neighbours in that class.

By induction on k, we shall obtain the desired family $\{U_\xi : \xi \in D\}$.

(c) Suppose that $x \in \Omega$. If x does not belong to any U_ξ, define $f(x) := 1$. Otherwise, define $f(x) := \inf\{\xi \in D : x \in U_\xi\}$.

If $x \in \mathrm{cl}(U_\alpha)$, where $D \ni \alpha < 1$, then, for any $\lambda \in D$ with $\lambda > \alpha$, $x \in U_\lambda$ by (16). Hence, $f(x) \le \lambda$, for any such λ. But $\inf\{\lambda \in D : \lambda > \alpha\} = \alpha$. We deduce that $f(x) \le \alpha$. (17)

If $x \in E$, then $x \in U_0$, and so $f(x) = 0$. If $x \in F$, then x does not belong to U_1, or, consequently, to any U_ξ, and so $f(x) = 1$.

(d) We finish by showing that f is continuous. By 5.4*(b)*, it makes no difference whether f is continuous as a mapping into $\mathbb{R}$ or as a mapping into I.

Given x_0, let $M \in \mathfrak{N}_{\mathbb{R}}(f(x_0))$. Assume, for convenience, that $0 < f(x_0) < 1$; the arguments if $f(x_0) = 0$ or $f(x_0) = 1$ are simpler, each being the appropriate half of the two-sided proof that follows.

Take dyadic numbers α, α' and β, β' such that

$$0 \le \alpha' < \alpha < f(x_0) < \beta < \beta' \le 1 \quad \text{and} \quad (a', \beta') \subseteq M .$$

If $x \in U_\beta$, then $f(x) \le \beta < \beta'$ by definition. (18)

Also, $x_0 \in U_\beta$. Indeed, if $x_0 \notin U_\beta$, then $x_0 \notin U_\xi$ for any $\xi < \beta$ by (16). Thus β is a lower bound for $\{\xi \in D : x \in U_\xi\}$, and, by the definition of f, $f(x_0) \geq \beta$; which is impossible.

(17) shows that, if $x \in \mathrm{cl}(U_\alpha)$, then $f(x) \leq \alpha$. It follows that $x_0 \notin \mathrm{cl}(U_\alpha)$; $\Omega \setminus \mathrm{cl}(U_\alpha)$ is an open set containing x_0. Further, if $y \notin \mathrm{cl}(U_\alpha)$, then $y \notin U_\xi$ for any $\xi \leq \alpha$, so $f(y) \geq \alpha$ as in the previous paragraph; hence $f(y) > \alpha'$. (19)

Putting together (18) and (19), $Y := U_\beta \setminus \mathrm{cl}(U_\alpha)$ is an open set, $x_0 \in Y$, and $f(Y) \subseteq (\alpha', \beta') \subseteq M$. Thus, f is continuous at x_0. □

The conclusion may be summarized informally as "if closed sets may be separated by open sets, they may also be separated by continuous functions".

The conception of this proof is rather easy, although the details are a little laborious. As I have already stressed, the point is the construction of a continuous function from entirely "abstract" data. In a metric space we already have continuous functions $d(x, A)$ and $d(x, B)$, and f can be defined by

$$f(x) := \frac{d(x, A)}{d(x, A) + d(x, B)} .$$

From the bare statement of 7.16, you might suppose that $A = f^{-1}(\{0\})$ and $B = f^{-1}(\{1\})$. In the proof, however, one sees clearly that this is not what has been asserted. Indeed, there may be essential hindrances to its being true.

Definition 7.17. Let Ω be a topological space. A subset A of Ω is a $\mathcal{G}_\delta$-set if it may be expressed as a countable intersection of open sets: $A = \bigcap_{n=1}^\infty U_n$, where each $U_n \in \mathcal{G}$. The class of all such sets is denoted $\mathcal{G}_\delta$. Complementarily, $B \in \mathcal{P}(\Omega)$ is an $\mathcal{F}_\sigma$-set if it is a countable union of closed sets: $B = \bigcup_{n=1}^\infty F_n$, where each $F_n \in \mathcal{F}$; and the class of such sets is called $\mathcal{F}_\sigma$.

Lemma 7.18. *In a metric space (Ω, d), every open set is $\mathcal{F}_\sigma$ and every closed set is $\mathcal{G}_\delta$: that is, $\mathcal{G} \subseteq \mathcal{F}_\sigma$ and $\mathcal{F} \subseteq \mathcal{G}_\delta$.*

Proof. The two statements are equivalent (by taking complements), so I prove the first. Let U be an open set. Define for $n \in \mathbb{N}$

$$F_n := \{x \in \Omega : B(x; 1/n) \subseteq U\} = \{x \in \Omega : d(x, \Omega \setminus U) \geq 1/n\} \subseteq U .$$

By 7.14*(b)* and 3.7*(a)*, this set is closed. But, by definition, $x \in U$ if and only if there is some $\epsilon > 0$ such that $B(x; \epsilon) \subseteq U$. Take n to be so large that $1/n \leq \epsilon$, and then $x \in F_n$. So $U = \bigcup_{n=1}^\infty F_n$, and is $\mathcal{F}_\sigma$.

$\mathcal{F}_\sigma$ is usually bigger than $\mathcal{G}$, since it contains closed sets as well as open ones.

Lemma 7.19. *Let Ω and Ψ be topological spaces, and $f : \Omega \longrightarrow \Psi$ a continuous mapping. If E is a $\mathcal{G}_\delta$ [or $\mathcal{F}_\sigma$] set in Ψ, then $f^{-1}(E)$ is $\mathcal{G}_\delta$ [or $\mathcal{F}_\sigma$] in Ω.* □

Corollary 7.20. *If $f : \Omega \longrightarrow \mathbb{R}$ is continuous (Ω being a topological space), then $f^{-1}(\{0\})$ and $f^{-1}(\{1\})$ are closed $\mathcal{G}_\delta$ in Ω.* □

One often writes $f^{-1}(0)$ instead of the more correct notation $f^{-1}(\{0\})$, and so on, just as one says "points are closed" rather than "singletons are closed".

Proposition 7.21. *Let Ω be a normal topological space, and suppose that A and B are disjoint nonnull closed subsets in Ω, and that A is $\mathcal{G}_\delta$. Then there is a continuous function $f : \Omega \longrightarrow I$ such that $f^{-1}(0) = A$ and $f^{-1}(1) \supseteq B$. If B is also $\mathcal{G}_\delta$, then f may be constructed so that $f^{-1}(1) = B$ and $f^{-1}(0) = A$.*

Proof. Suppose that $A = \bigcap_{n=1}^{\infty} G_n$, where each G_n is open; take $H_n := G_n \setminus B$, and $A = \bigcap_{n=1}^{\infty} H_n$, where H_n is open and $B \subseteq \Omega \setminus H_n$. By Urysohn's lemma, there is a continuous function $f_n : \Omega \longrightarrow [0,1]$ such that $f(A) = \{0\}$ and $f(\Omega \setminus H_n) = \{1\}$. For each $x \in \Omega$, define $F(x) := \sum_{n=1}^{\infty} 2^{-n} f_n(x)$. This sum converges uniformly, by comparison with $\sum 2^{-n}$. It therefore defines a continuous function with values in I, and it is clear that its value on B is 1, whilst, for any $x \notin A$, there is some n such that $x \notin H_n$, so that $F(x) \geq 2^{-n}$. ****** □

I have assumed some facts here about uniform convergence, and perhaps I should say what they are, since they will soon be used again. Suppose that $(F_d)_{d \in D}$ is a net (commonly a sequence) of continuous functions $\Omega \longrightarrow \mathbb{R}$. [The example I have in mind above is that $F_m = \sum_{n=1}^{m} 2^{-n} f_n$ in the usual sense of addition of functions, i.e. pointwise addition: $(\forall x \in \Omega)\ \ F_m(x) := \sum_{n=1}^{m} 2^{-n} f_n(x)$. As this is a finite sum of continuous functions, it is continuous.] We say that the net converges *uniformly on* Ω to the function F if

$$(\forall \epsilon > 0)(\exists d_0 \in D)(\forall x \in \Omega) \quad d \geq d_0 \Longrightarrow |F_n(x) - F(x)| < \epsilon\,.$$

This is the usual definition of convergence for each individual $x \in \Omega$ of the numerical net $(F_d(x))_{d \in D}$ to the limit $F(x)$, but with the added demand that the d_0 whose existence (for the given ϵ) is asserted should "work" for the given ϵ and *for all x simultaneously*.

The standard example (apologies to those who have done 312) is the sequence of functions $F_n : I \longrightarrow I$, where $F_n(t) := t^n$ for each n. In this case the numerical sequences (t^n) all converge; the limit is 1 when $t = 1$ and is 0 otherwise. If you take for instance $\epsilon := \frac{1}{2}$, there is no N such that $|t^N| < \frac{1}{2}$ for $0 \leq t < 1$. Whatever N is chosen, $t_N \geq \frac{1}{2}$ for some t very close to 1. So the convergence of this sequence of functions is non-uniform. [On the other hand, in the proof above, where $F_m := \sum_{n=1}^{m} 2^{-n} f_n$, the sequence does converge uniformly to F, since

$$\begin{aligned}(\forall x \in \Omega) \quad |F(x) - F_m(x)| &= \left|\sum_{n=m+1}^{\infty} 2^{-n} f_n(x)\right| \\ &\leq \sum_{n=m+1}^{\infty} 2^{-n} = 2^{-m}\,.]\end{aligned}$$

The important result about uniform convergence is the following.

Lemma 7.22. *Let $(F_d : \Omega \longrightarrow \mathbb{R})_{d \in D}$ be a net of functions on the topological space Ω, and suppose it converges* uniformly *to a limit function $F : \Omega \longrightarrow \mathbb{R}$. If each F_d is continuous at $x_0 \in \Omega$, then F is also continuous at x_0.*

Proof. Take any $\epsilon > 0$. There exists some $d_0 \in D$ such that

$$(\forall x \in \Omega) \quad d \geq d_0 \Longrightarrow |F_d(x) - F(x)| < \tfrac{1}{3}\epsilon\,. \tag{20}$$

But F_{d_0} is continuous at x_0, so there exists a neighbourhood M of x_0 with

$$x \in M \Longrightarrow |F_{d_0}(x) - F_{d_0}(x_0)| < \tfrac{1}{3}\epsilon\,. \tag{21}$$

Thus, if $x \in M$,

$$\begin{aligned} |F(x) - F(x_0)| &\leq |F(x) - F_{d_0}(x)| + |F_{d_0}(x) - F_{d_0}(x_0)| + |F_{d_0}(x_0) - F(x_0)| \\ &< \tfrac{1}{3}\epsilon + \tfrac{1}{3}\epsilon + \tfrac{1}{3}\epsilon \qquad \text{by (20) and (21)} \\ &= \epsilon\,. \end{aligned}$$

Thus, F is continuous at x_0. □

It follows that, indeed, the function F constructed in the proof of 7.21 is continuous, as we required it to be.

Definition 7.23. The topological space Ω is *perfectly normal* or T_6 if it is normal and if every closed subset of Ω is $\mathcal{G}_\delta$.

It requires no new argument to observe that a metric space is perfectly normal.

Urysohn's lemma is a special case of

Theorem 7.24. (The Tietze-Urysohn extension theorem.) *Let Ω be a normal space, and let A be a closed subset. If A is given the subspace topology and $f : A \longrightarrow \mathbb{R}$ is continuous, then there is a continuous function $g : \Omega \longrightarrow \mathbb{R}$ such that $g|A = f$. Furthermore, if f is bounded above (or below, or both), then g may be chosen to have the same bounds; and if a bound for f is not attained, it is not attained by g.*

The last sentence means that, if $f(x) \leq \alpha$ for all $x \in A$, then g may be constructed so that $g(x) \leq \alpha$ for all $x \in \Omega$; if $f(x) < \alpha$ for all $x \in A$, then $g(x) < \alpha$ for all $x \in \Omega$; and similarly for lower bounds.

Proof. Suppose to begin with that $\phi : A \longrightarrow [-c, c]$ is continuous, for some $c > 0$. The sets $A_- := \phi^{-1}([-c, -\frac{1}{3}c])$ and $A_+ := \phi^{-1}([\frac{1}{3}c, c])$ are disjoint and closed in A, and consequently are closed and disjoint in Ω. By Urysohn's lemma 7.16, there is a continuous $k : \Omega \longrightarrow [0,1]$ such that $k(A_1) = \{0\}$ and $k(A_+) = \{1\}$. Define

$$\psi(x) := \tfrac{2}{3}c(k(x) - \tfrac{1}{2})\,;$$

then $|\psi(x)| \leq \frac{1}{3}c$ for all $x \in \Omega$, $|\phi(x) - \psi(x)| \leq \frac{2}{3}c$ for all $x \in A$. [Check this].

So now suppose that $f : A \longrightarrow [-1, 1]$ is continuous. By the previous paragraph (taking $\phi := f$, $c := 1$), there is a continuous function $g_1 : \Omega \longrightarrow \mathbb{R}$ such that $|g_1(x)| \leq \frac{1}{3}$ for all $x \in \Omega$ and $|f(x) - g_1(x)| \leq \frac{2}{3}$ for all $x \in A$. This is the first step of an induction; suppose now that $g_1, g_2, \ldots, g_k : \Omega \longrightarrow \mathbb{R}$ have been defined so that, for $1 \leq i \leq k$ and for all $x \in \Omega$,

$$|g_i(x)| \leq \tfrac{1}{3}\left(\tfrac{2}{3}\right)^{i-1}, \tag{22}$$

and, for all $x \in A$,

$$\left|f(x)-\sum_{i=1}^{k} g_i(x)\right| \le \left(\tfrac{2}{3}\right)^k . \tag{23}$$

Take $\phi := f - \sum_{i=1}^{k} g_i$ and $c := \left(\frac{2}{3}\right)^n$ above, to deduce the existence of a continuous function $g_{k+1} : \Omega \longrightarrow \mathbb{R}$ such that $|g_{k+1}(x)| \le \frac{1}{3}\left(\frac{2}{3}\right)^k$ and

$$(\forall x \in A) \quad \left|f(x)-\sum_{i=1}^{k} g_i(x) - g_{k+1}(x)\right| \le \left(\tfrac{2}{3}\right)^{k+1} .$$

In this way we obtain a series of functions, $\sum g_i$, satisfying (22) and (23) for all i. On Ω it converges uniformly, by comparison with the numerical series $\sum \frac{1}{3}\left(\frac{2}{3}\right)^{i-1}$, which converges with sum 1; thus, certainly, $|\sum_{i=1}^{\infty} g_i(x)| \le 1$ for all $x \in \Omega$. Define $g(x) := \sum_{i=1}^{\infty} f(x)$ for each $x \in \Omega$; then g is continuous by 7.22, and $g|A = f$ by (23). And $|g(x)| \le 1$ for all $x \in \Omega$, so $g : \Omega \longrightarrow [-1,1]$.

Of course if f attains the value 1 at some point of A, then g must necessarily do so as well. But suppose that $f(x) < 1$ for all $x \in A$. It is at least possible that our construction allows g to take the value 1 at some points of Ω. Define

$$B := g^{-1}(1)\,;$$

this is a closed set in Ω, and is disjoint from A (by our assumption on f). By 7.16, there is a continuous function $\chi : \Omega \longrightarrow [0,1]$ such that $\chi(B) = \{0\}$ and $\chi(A) = \{1\}$. The function $h(x) := g(x)\chi(x) : \Omega \longrightarrow [\,-1,1]$ is continuous and extends f, but does not take the value 1. A similar trick can be used to construct extensions of f that do not take the value -1 if f does not; or either of the values $-1, 1$ if f does not.

If $-\infty < a < b < \infty$ and $f_0 : A \longrightarrow [a,b]$ is continuous, then set

$$(\forall x \in A) \quad f(x) := \frac{2f_0(x) - (a+b)}{b-a} .$$

f is a continuous function $A \longrightarrow [-1,1]$, and does not take the value 1 (or -1) when f_0 does not take the value b (or a). It has a continuous extension to $g : \Omega \longrightarrow [-1,1]$, which does not take the value 1 or -1 if f does not; and then

$$g_0(x) := \tfrac{1}{2}\{(b-a)g(x) + (a+b)\}$$

is an extension of f_0 to a continuous function $g_0 : \Omega \longrightarrow [a,b]$, which does not take the value a (or b) if f_0 does not.

Hence, if $f_0 : A \longrightarrow \mathbb{R}$ is a bounded continuous function and $a := \inf f_0(A)$, $b := \sup f_0(A)$, there is a continuous extension of f_0 to a bounded function $g_0 : \Omega \longrightarrow \mathbb{R}$ with the same supremum and infimum, and it attains either of them on Ω if and only if f_0 attains it originally on A.

There is a similar argument if f_0 is bounded below, $a := \inf f_0(A)$, but not above; then one defines for $x \in A$

$$f(x) := \frac{f_0(x) - (a+1)}{1 + f_0(x) - a} ,$$

which maps into $[-1,1)$. f may be extended to $g : \Omega \longrightarrow A$, as before, which does not attain the value 1 on Ω and attains the value -1 only if f attains it on A. Define

$$(\forall x \in \Omega) \quad g_0(x) := a - 1 + \frac{2}{1 - g(x)} \ ;$$

this is the desired extension of f_0, which attains its infimum a at a point of Ω if and only if f_0 attains its infimum on A.

Finally, if $f_0 : A \longrightarrow \mathbb{R}$ is bounded neither above nor below, define

$$(\forall x \in A) \quad f(x) := \frac{f_0(x)}{1 + |f_0(x)|} \ .$$

Then $f : A \longrightarrow \mathbb{R}$ has unattained infimum -1 and supremum $+1$. It has a continuous extension $g : \Omega \longrightarrow (-1, 1)$ (i.e. g too does not attain its supremum or infimum), and $g_0 : \Omega \longrightarrow \mathbb{R} : x \mapsto g(x)/(1 - |g(x)|)$ is a continuous extension of f_0.☐

Remark 7.25. Although the above proof is long, that is partly because there are several distinct statements to deal with.

There are two other versions of this theorem that I know of. Tietze's original theorem, proved in 1915, was for metric spaces. As for Urysohn's lemma, the metric can be used to construct the extension, at least when f is bounded, by a simple formula. One may suppose that $\inf\{f(x) : x \in A\} = 1$ and $\sup\{f(x) : x \in A\} = 2$, and then one can set

$$(\forall x \in \Omega) \quad g(x) := \begin{cases} \inf \left\{ \dfrac{f(y) d(x, y)}{d(x, A)} : y \in A \right\} & \text{when } x \notin A \text{, and} \\ f(x) & \text{when } x \in A \, . \end{cases}$$

The proof that g is continuous on Ω and extends f is not entirely trivial.

There is also a famous theorem due to Dugundji about the extension of continuous functions with values in a locally convex topological vector space. But it requires paracompactness of Ω, which we have not yet discussed.

The general question of extension of continuous mappings with values in a topological space led Borsuk to invent the "theory of retracts".

If one studies the proof of Urysohn's lemma, it becomes apparent that normality is used in an essential way; even if both A and B are singletons, the sets V_α will not usually be. Thus, in particular, there is no analogous result about "using a continuous function to separate a point x from a closed set F such that $x \notin F$" in a regular space, despite the similarity of the definitions of T_3 and T_4. However, a "functional separation property" for points and closed sets has some importance later.

Definition 7.26. The topological space Ω is *Tikhonov* [the spelling is variable, for obvious reasons], or *completely regular*, or — this is no longer a joke, despite presumably being invented as one — $T_{3\frac{1}{2}}$, if it is T_1 and, for any $x \in \Omega$ and any closed subset F of Ω such that $x \notin F$, there is a continuous function $f : \Omega \longrightarrow [0, 1]$ such that $f(x) = 0$ and $f(F) = \{1\}$.

An equivalent formulation is that, for each point $x \in \Omega$, the sets

$$\{y \in \Omega : |f(y) - f(x)| < \epsilon\}$$

where f ranges over the continuous functions $\Omega \longrightarrow \mathbb{R}$ and ϵ over the positive real numbers, form a base of neighbourhoods for x. T_4 implies $T_{3\frac{1}{2}}$ by Urysohn's lemma, and $T_{3\frac{1}{2}}$ implies T_3.

8 Compactness.

It is quite common to give the definitions and so on for *spaces*; however, most applications deal with compact *subsets* of a space, so I shall begin from that.

Definition 8.1. Let A be a subset of a set Ω. A *covering* of A is a class $\mathcal{U}$ of subsets of Ω (that is, $\mathcal{U} \subseteq \mathcal{P}(\Omega)$) such that $\bigcup_{U \in \mathcal{U}} U \supseteq A$. If every $U \in \mathcal{U}$ is an open set in Ω, $\mathcal{U}$ is described as an *open covering* of A. Similarly, if every $U \in \mathcal{U}$ is a closed set in Ω, $\mathcal{U}$ is a *closed covering* of A.

Let $\mathcal{U}_1$ and $\mathcal{U}_2$ be coverings of A. $\mathcal{U}_1$ is a *subcovering* of $\mathcal{U}_2$ (sc. "of A") if $\mathcal{U}_1 \subseteq \mathcal{U}_2$ as subsets of $\mathcal{P}(\Omega)$; that is, if every member of $\mathcal{U}_1$ is a member of $\mathcal{U}_2$. $\mathcal{U}_1$ is a *refinement* of $\mathcal{U}_2$ if every member of $\mathcal{U}_1$ is *included in* a member of $\mathcal{U}_2$:

$$(\forall U \in \mathcal{U}_1)(\exists V \in \mathcal{U}_2) \quad U \subseteq V \, .$$

Notice that, whereas a subcovering may have fewer members than the original covering, a refinement may have many more.

Definition 8.2. The subset A of the topological space Ω is *compact* if every open covering of A has a finite subcovering.

The definition is sometimes formulated in two other ways: every open covering of A has a finite open refinement, or has a finite refinement. It is clear that these statements are equivalent to 8.2, and their only advantage over it is that they are closer to some later definitions.

As I remarked, compactness is often defined for "whole spaces", and then a compact *subset* is a subset that is compact as a topological space with respect to the subspace topology. This is equivalent to my definition, and emphasizes that compactness is "intrinsic to the set A" — it depends only on the subspace topology of A, not on any subtler properties of the way A lies in Ω.

My definition is pretty standard, but in some sources you will find slightly different formulations. Engelking demands that a compact space be T_2; what I call "compact", he calls "quasi-compact". I am not sure why. Russian authors, by contrast, used to use the word "bicompact" to mean "compact T_2".

Compactness is a remarkable property. If A is a finite set, it is evidently compact (why?), but except for that trivial case it is at first sight difficult to imagine other examples. Of course, I gave some in 312 and 441.

Ω will denote a topological space from now on.

Lemma 8.3. *Let K be a compact subset of Ω and C a closed subset of Ω. Then $K \cap C$ is compact.* [A (relatively) closed subset of a compact set is compact.]

Proof. Let $\mathcal{U}$ be an open covering of $K \cap C$. Then $\mathcal{U} \cup \{\Omega \setminus C\}$ is an open covering of K. Hence, it has a finite subcovering $\mathcal{U}'$. Define $\mathcal{U}'' := \mathcal{U}' \setminus \{\Omega \setminus C\}$; this is an open covering of $K \cap C$, because any point of $K \cap C$ belongs to some member of $\mathcal{U}'$ but that member cannot be $\Omega \setminus C$; and it is a subcovering of $\mathcal{U}$, and finite. □

Lemma 8.4. *Let Ω be Hausdorff. Then any compact subset K of Ω must be closed. Furthermore, if C is a second compact subset disjoint from K, there are open subsets U, V such that $U \cap V = \emptyset$, $C \subseteq U$, and $K \subseteq V$.*

Proof. Suppose $x \notin K$. Then, for any $k \in K$, there are open sets $U(k), V(k)$ such that $x \in U(k)$, $k \in V(k)$, and $U(k) \cap V(k) = \emptyset$. Thus $\{V(k) : k \in K\}$ is an open covering of K, so it has a finite subcovering $\{V(k_i) : 1 \leq i \leq m\}$. Let

$$\widehat{U}(x) := \bigcap_{i=1}^{m} U(k_i)\,.$$

Then $\widehat{U}(x)$ is an open set; it contains x (as each $U(k_i)$ does). Set $\widehat{V}(x) := \bigcup_{i=1}^{m} V(k_i)$, and then

$$\begin{aligned}\widehat{U}(x) \cap K \subseteq \widehat{U}(x) \cap \widehat{V}(x) &= \bigcup_{i=1}^{m} (\widehat{U}(x) \cap V(k_i)) \\ &\subseteq \bigcup_{i=1}^{m} (U(k_i) \cap V(k_i)) = \emptyset\,.\end{aligned}$$

Hence, in the first place, $\widehat{U}(x) \subseteq \Omega \setminus K$; this proves that $\Omega \setminus K$ is a neighbourhood of x; and, as x was an arbitrary point of $\Omega \setminus K$, $\Omega \setminus K$ is open in Ω.

Secondly, applying the argument to each $x \in C$ gives open sets $\widehat{U}(x), \widehat{V}(x)$ such that $\widehat{U}(x) \cap \widehat{V}(x) = \emptyset$, $x \in \widehat{U}(x)$, $K \subseteq \widehat{V}(x)$. Now $\{\widehat{U}(x) : x \in C\}$ is an open covering of C, so has a finite subcovering $\{\widehat{U}(x_k) : 1 \leq k \leq n\}$. Let

$$U := \bigcup_{k=1}^{n} \widehat{U}(x_k)\,, \quad V := \bigcap_{k=1}^{n} \widehat{V}(x_k)\,.$$ □

This result need not be true if Ω is not T_2. Similarly,

Lemma 8.5. *Suppose that Ω is compact Hausdorff. Then it is normal.*

Proof. If A, B are disjoint closed sets in Ω, they are compact by 8.3, so 8.4 applies. □

Regular****

Lemma 8.6. *Let K, L be compact subsets of Ω. Then $K \cup L$ is compact.*

Proof. If $\mathcal{U}$ is an open covering of $K \cup L$, then it is an open covering of K, in which capacity it has a finite subcovering $\mathcal{U}_1$, and also an open covering of L, having a finite subcovering $\mathcal{U}_2$. Then $\mathcal{U}_1 \cup \mathcal{U}_2$ is a finite subcovering of $K \cup L$, since

$$\bigcup_{U \in \mathcal{U}_1 \cup \mathcal{U}_2} = \Big(\bigcup_{U \in \mathcal{U}_1}\Big) \cup \Big(\bigcup_{U \in \mathcal{U}_2} U\Big) \supseteq K \cup L\,.$$ □

Lemma 8.7. *The subset A of Ω is compact if and only if it has the "finite intersection property": if $\mathcal{C}$ is a class of closed sets such that, for any finite subclass $\{C_1, C_2, \ldots, C_k\}$ of $\mathcal{C}$, $A \cap C_1 \cap C_2 \cap \cdots \cap C_k \neq \emptyset$, then $\bigcap_{C \in \mathcal{C}} (A \cap C) \neq \emptyset$.*

Proof. Complementation (with respect to K) and the contrapositive. □

Possibly the most striking fact about compact sets (though it is not profound) is the following:

Lemma 8.8. *If $f : \Omega \longrightarrow \Psi$ is continuous and A is compact in Ω, then $f(A)$ is compact in Ψ.*

Proof. Let $\mathcal{U}$ be an open covering of $f(A)$. Then $\{f^{-1}(U) : U \in \mathcal{U}\}$ is an open covering of A. As such, it has a finite subcovering, $\{f^{-1}(U_k) : 1 \leq k \leq n\}$, and then $\{U_k : 1 \leq k \leq n\}$ is a finite subcovering of $f(A)$. □

The last step is not quite obvious, the point being that, for any $B \subseteq \Psi$, $f(A \cap f^{-1}(B)) = f(A) \cap B$.

Corollary 8.9. *Suppose that Ω is compact, Ψ is Hausdorff, and $f : \Omega \longrightarrow \Psi$ is continuous. Then f is closed (recall* 3.13(b)*); that is, for every closed set A in Ω, $f(A)$ is closed in Ψ.*

Proof. 8.3, 8.8, and 8.4. □

This has a remarkable consequence:

Corollary 8.10. *If f in* 8.9 *is a bijection, then it is a homeomorphism. If f is a surjection, then the topology on Ψ is the topology induced by f from the topology of Ω (see* 5.20*), i.e. the quotient topology* (5.25). **** □

This is interesting because in the familiar *algebraic* categories (groups, rings, fields, vector spaces) — although not in some less familiar ones like semigroups — a bijective homomorphism is automatically an isomorphism, and, in the category of topological spaces and continuous maps, the corresponding statement is certainly false (as the discrete topology shows). But in the category of compact Hausdorff spaces and continuous maps, it is again true.

Lemma 8.11. *The subset A is compact if and only if every net in A has a subnet convergent to a point of A.*

This is, of course, the immediate analogue of "sequential compactness" in the case of metric spaces. Before proving it, we need a little introduction.

Lemma 8.12. *x is a cluster point of $(x_d)_{d \in D}$ if and only if it is the limit of a convergent subnet.*

Proof. Let x be a cluster point. Define $E := \mathfrak{N}(x) \times D$, with partial order

$$(M_1, d_1) \leq (M_2, d_2) \Longleftrightarrow (M_1 \supseteq M_2 \;\;\&\;\; d_1 \leq d_2)\,.$$

Given $e = (M, d) \in E$, *choose* $\phi(e) \in D$ to be such that $\phi(e) \geq d$ and $x_{\phi(e)} \in M$. This is possible, by the definition 10.2. The mapping ϕ thus defined satisfies Definition 6.12, and so is a subnet. But it also converges to x (6.9).

The converse argument is left as an exercise. □

Proof of 8.11. Suppose that A is compact, and that $(x_d)_{d\in D}$ is a net in A. If it has no cluster points in A, then, for each $x \in A$, there is some open set $U(x)$ with $x \in U(x)$ and $\{d \in D : x_d \in U(x)\}$ not cofinal in D; this means that there is some $d(x) \in D$ such that, if $d(x) \leq d$, $x_d \notin U(x)$. Now $\{U(x) : x \in A\}$ is an open covering of A, so it has a finite subcovering $\{U(x_i) : 1 \leq i \leq m\}$.

As D is directed, there is some d_0 such that $d(x_i) \leq d_0$ for $1 \leq i \leq m$. Suppose that $d_0 \leq d \in D$; then $x_d \notin U(x_i)$ for $1 \leq i \leq m$. And this is absurd, since it implies that $x_d \notin A$. This proves that $(x_d)_{d\in D}$ has a cluster point in A.

Suppose A is not compact. By 8.7, there is a class $\mathcal{C}$ of closed sets the intersection of all of which does *not* meet A, but such that any finite intersection of them *does* meet A. Direct the family D of finite subclasses of $\mathcal{C}$ by inclusion. The result is a directed set. For each $d \in D$, choose an element $x_d \in A \cap (\bigcap_{C\in d} C)$. This defines a net in A. I claim that it has no cluster point in A.

Indeed, if x were a cluster point, then, for any $M \in \mathfrak{N}(x)$, $\{d \in D : x_d \in M\}$ would be cofinal. This implies x is adherent to each $C \in \mathcal{C}$, and so belongs to each C. And this in turn contradicts the hypothesis that $A \cap (\bigcap_{C\in\mathcal{C}} C) = \emptyset$. □

Although the above argument is satisfactory enough, it uses the axiom of choice repeatedly in a rather irresponsible way (and needs the unpleasant definition of a subnet). Filters make for a better theory, albeit one based on the same ideas.

Definition 8.13. Let $\mathfrak{F}$ be a filter (or filter base) in Ω. $x \in \Omega$ is a *cluster point* of $\mathfrak{F}$ if, for any $M \in \mathfrak{N}(x)$ and $F \in \mathfrak{F}$, $M \cap F \neq \emptyset$; or equivalently $x \in \bigcap_{F\in\mathfrak{F}} \mathrm{cl}(F)$. [Of course $\{\mathrm{cl}\, F : F \in \mathfrak{F}\}$ is a filter in the class $\mathcal{F}$ of closed sets in Ω.]

Lemma 8.14. *x is a cluster point of $\mathfrak{F}$ if and only if it is a limit of a filter refining $\mathfrak{F}$.*

Proof. If x is a cluster point of $\mathfrak{F}$, then $\{M \cap F : M \in \mathfrak{N}(x)\ \&\ F \in \mathfrak{F}\}$ is a filter base in Ω, and the filter it generates is clearly a refinement of $\mathfrak{F}$ that converges to x. Conversely, if $\mathfrak{F} \subseteq \mathfrak{F}_1 \to x$, then $\mathfrak{F} \subseteq \mathfrak{F}_1 \supseteq \mathfrak{N}(x)$. If $F \in \mathfrak{F}$ and $M \in \mathfrak{N}(x)$, then $F, M \in \mathfrak{F}_1$, and so $F \cap M \neq \emptyset$. □

Lemma 8.15. *A is compact if and only if every filter in A has a refinement convergent in A, or equivalently has a cluster point in A.*

Proof. Suppose A is compact, but the filter $\mathfrak{F}$ in A lacks a cluster point in A. Then, for each $x \in A$, there is an open set $U(x)$ such that $x \in U$ and, for some $F(x) \in \mathfrak{F}$, $U(x) \cap F(x) = \emptyset$. Then $\{U(x) : x \in A\}$ is an open covering of A, with a finite subcovering $\{U(x_i) : 1 \leq i \leq m\}$. But $\bigcap_{i=1}^m F(x_i) \in \mathfrak{F}$ (as $\mathfrak{F}$ is a filter), and

$$\bigcap_{j=1}^m F(x_j) = A \cap \left(\bigcap_{j=1}^m F(x_j)\right) \subseteq \left(\bigcup_{i=1}^m U(x_i)\right) \cap \left(\bigcap_{j=1}^m F(x_j)\right)$$
$$\subseteq \bigcup_{i=1}^m (U(x_i) \cap F(x_i)) = \emptyset\,.$$

This is absurd, as $\mathfrak{F}$ is a filter in A. So $\mathfrak{F}$ must have a cluster point in A.

Conversely, suppose that A is not compact. By 8.7, there is a class $\mathcal{C}$ of closed sets of Ω such that any finite intersection of them meets A, but the intersection of them all does not. Thus $\{A \cap C : C \in \mathcal{C}\}$ is a filter base in A. But, by construction, it has no cluster point in A. □

These proofs do not require choices, unlike the ones with nets. However, the mention of "refinements" of the filters suggests something else. Let Φ be the class of all filters in A. Then "refinement" is a partial order in Φ; $\mathfrak{F}_1 \leq \mathfrak{F}_2$ means "$\mathfrak{F}_2 \supseteq \mathfrak{F}_1$". It is easily checked (with a little argument *****) that this partial order is inductive (§4**C**), and so, by Zorn's lemma, any filter is included in a maximal filter.

Definition 8.16. A maximal filter (in the class of filters of subsets of the set A) is an *ultrafilter* in A. A base for an ultrafilter may be called an *ultrabase*. (These concepts are purely set-theoretical. Notice that an ultrabase is not necessarily maximal itself in any obvious class.)

If one considers a topology on A and filters in the class of closed subsets of A, one can similarly construct ultrafilters; they are sometimes called *maximal closed bases*. The name is consistent, because a filter base (in whatever class of subsets) that is maximal in the class of filter bases will itself be a filter.

An ultrafilter or maximal closed base $\mathfrak{F}$ is *free* if $\bigcap_{F\in\mathfrak{F}} F = \emptyset$; *finite* intersections of its members are nonempty by definition. It is *principal* if there exists some $a \in A$ such that $\mathfrak{F}$ consists of all the sets [or all the closed sets, when $\mathfrak{F}$ is a maximal closed base] containing a.

Lemma 8.17. (a) *An ultrafilter is free if and only if it is not principal.*

(b) *If A is* T_1 *in the relative topology, a maximal closed base in A is free if and only if it is not principal.* □

Lemma 8.18. *A filter base $\mathfrak{B}$ in a set A is an ultrabase if and only if any of the following equivalent conditions holds.*

(a) *If $E \in \mathcal{P}(A)$ and $B \cap E \neq \emptyset$ for all $B \in \mathfrak{B}$, then there is some $B_0 \in \mathfrak{B}$ such that $B_0 \subseteq E$. (Then E belongs to the filter generated by $\mathfrak{B}$.)*

(b) *If $E \in \mathcal{P}(A)$, then either there is some $B_0 \in \mathfrak{B}$ such that $B_0 \subseteq E$ or there is some $B_1 \in \mathfrak{B}$ such that $B_1 \subseteq A \setminus E$.*

(c) *If $E_1, E_2, \ldots, E_k \in \mathcal{P}(A)$ and $E_1 \cup E_2 \cup \cdots \cup E_k \supseteq B_0 \in \mathfrak{B}$, then there exists an index i, with $1 \leq i \leq k$, such that $E_i \supseteq B_1 \in \mathfrak{B}$ for some $B_1 \in \mathfrak{B}$.*

(d) *If $E_1, E_2, \ldots, E_k \in \mathcal{P}(A)$, $B_0 \in \mathfrak{B}$, and $E_1 \cap E_2 \cap \cdots \cap E_k \cap B_0 = \emptyset$, then there exists an index i, with $1 \leq i \leq k$, such that $E_i \cap B_1 = \emptyset$ for some $B_1 \in \mathfrak{B}$.* □

Of course the statements for filters and ultrafilters are more elegant.

Lemma 8.19. *The set A in the topological space Ω is compact if and only if every ultrafilter in A converges (or, equivalently, every ultrabase converges).* $\square$

Lemma 8.20. *Let $\mathfrak{B}$ be an ultrabase in the set Ω, and let $f : \Omega \longrightarrow \Psi$ be a mapping. Then $f(\mathfrak{B})$ is an ultrabase in Ψ.*

Proof. $f(\mathfrak{B})$ is a filter base by 6.20. If it is not an ultrabase, then, by 8.18*(a)*, there is some $E \in \mathcal{P}(\Omega)$ such that $E \cap f(B) \neq \emptyset$ for every $B \in \mathfrak{B}$, but there is no $B_0 \in \mathfrak{B}$ such that $f(B_0) \subseteq E$. So $f^{-1}(E) \cap B \neq \emptyset$, but $B \not\subseteq f^{-1}(E)$, for every $B \in \mathfrak{B}$. This, however, contradicts 8.18*(a)* for $\mathfrak{B}$. So $f(\mathfrak{B})$ is an ultrabase. $\square$

Proposition 8.21. *Suppose that, for $\beta \in B$, Ω_β is a topological space; let $\Omega := \prod_{\beta \in B} \Omega_\beta$, and let $\pi_\beta : \Omega \longrightarrow \Omega_\beta$ be the coordinate projection. A filter base $\mathfrak{B}$ in Ω converges to $x \in \Omega$ if and only if, for each $\beta \in B$, the filter base $\pi_\beta(\mathfrak{B})$ in Ω_β converges to $\pi_\beta(x)$ in Ω_β.*

Proof. "Only if" results from 6.21. Suppose that, for each $\beta \in B$, $\pi_\beta(\mathfrak{B}) \to x_\beta$ in Ω_β, and let $x := (x_\beta)_{\beta \in B}$. A neighbourhood M of x in Ω will include a "basic open set" of the form $W := \pi_{\beta_1}^{-1}(U_1) \cap \pi_{\beta_2}^{-1}(U_2) \cap \cdots \cap \pi_{\beta_k}^{-1}(U_k)$ (see (10) of 5.15), where $x_{\beta_i} \in U_i$ for $1 \leq i \leq k$. But then there is some $B_i \in \mathfrak{B}$ such that $\pi_{\beta_i}(B_i) \subseteq U_i$, which is the same as $B_i \subseteq \pi_{\beta_i}^{-1}(U_i)$; hence $B_1 \cap B_2 \cap \cdots \cap B_k \subseteq W$, but, as $\mathfrak{B}$ is a filter base, there is some $B_0 \in \mathfrak{B}$ such that $B_0 \subseteq B_1 \cap B_2 \cap \cdots \cap B_k$, and so $B_0 \subseteq M$. This shows that $\mathfrak{B} \to x$ in Ω. $\square$

Now for the big theorem on compactness.

Theorem 8.22. (Tikhonov's theorem.) *Let the topological spaces Ω_β, for $\beta \in B$, be compact. Then their product $\Omega := \prod_{\beta \in B} \Omega_\beta$ is compact.*

Proof. Let $\mathfrak{F}$ be an ultrafilter [or ultrabase] in Ω. For each $\beta \in B$, $\pi_\beta(\mathfrak{F})$ is an ultrabase in Ω_β, by 8.20. By 8.19, $\pi_\beta(\mathfrak{F})$ converges in Ω_β, say to x_β. In Ω, by 8.21, $\mathfrak{F} \to (x_\beta)$. The result follows by 8.19. $\square$

This is the most memorable proof of the theorem that I know, because, as one hopes will ideally be the case, all the difficult bits have been moved into the "abstract nonsense" about filters. I believe the proof was first formulated in this way by Bourbaki, although previous proofs were essentially equivalent. In particular, the (first) proof given by Kelley relies on the "Alexander subbase lemma" instead, which was not originally formulated using filters. Here it is:

Lemma 8.23. *Let $\mathcal{S}$ be a subbase for the topology of the topological space Ω. Then Ω is compact if and only if every covering by members of $\mathcal{S}$ has a finite subcovering.*

Proof. "Only if" is obvious. Suppose, on the other hand, that every covering by members of $\mathcal{S}$ has a finite subcovering, but that Ω is not compact. Then there is an ultrafilter $\mathfrak{F}$ in Ω that is not convergent. For each $x \in \Omega$, there is some neighbourhood M of x that is not in $\mathfrak{F}$; by 8.18*(a)*, there exists $F_0 \in \mathfrak{F}$ such that $M \cap F_0 = \emptyset$. But M must include a "basic open neighbourhood" of x of the form $S_1 \cap S_2 \cap \cdots \cap S_k$, where each $S_i \in \mathcal{S}$, and then $S_1 \cap \cdots \cap S_k \cap F_0 = \emptyset$. By 8.18*(d)*, there is some $F_1 \in \mathfrak{F}$ and there is some index i, $1 \leq i \leq k$, such that $S_i \cap F_1 = \emptyset$.

This proves that, for each $x \in \Omega$, there is some $S(x) \in \mathcal{S}$ such that $x \in S(x)$ and, for some $F(x) \in \mathfrak{F}$, $S(x) \cap F(x) = \emptyset$. Thus, $\{S(x) : x \in \Omega\}$ is an open covering of Ω by sets of $\mathcal{S}$; by hypothesis, it has a finite subcovering $\{S(x_i) : 1 \leq i \leq m\}$, and then $F := \bigcap_{i=1}^{m} F(x_i) \in \mathfrak{F}$ and

$$F = F \cap \Omega = \bigcup_{i=1}^{m} (F \cap S(x_i)) = \emptyset .$$

This is absurd. So, in fact, Ω must be compact. □

In the case of finite Cartesian products, it is not necessary to use the axiom of choice at all; one may prove quite easily that the product of two compact spaces is compact****. But in general the axiom of choice is indispensable in the proof of Tikhonov's theorem. (It is, of course, needed to assure us that the product space is nonempty if each factor is nonempty; but it is conceivable that the theorem might remain true even if, in some cases, the product were empty.) In fact, Tikhonov's theorem is *equivalent* to the Multiplicative Axiom, and the implication Tikhonov $\Longrightarrow$ Multiplicative is a rather easy exercise.*****

It scarcely needs pointing out that so far I have given in this course no non-trivial examples of compact sets.

Proposition 8.24. *Let* $a, b \in \mathbb{R}$. *The interval* $[a, b]$ *is compact.*

This is usually quoted as "closed bounded intervals are compact". In 312 and 441, I present a well-known proof that quite explicitly relies on Dedekind's Axiom for $\mathbb{R}$ (by appealing to the properties of suprema). The Proposition is in a sense equivalent to Dedekind's Axiom, and to the General Principle of Convergence; that is, each can be deduced from either of the others, granted that $\mathbb{R}$ is an ordered field. This is why these three basic facts about $\mathbb{R}$ are expounded in different orders in various sources.

Proof. It is clear that the cases when $a \geq b$ are trivial, and that, for $a < b$, it suffices to consider the interval $[0, 1]$.

Let $\Omega := \prod_{k=1}^{\infty} \{0, 1\}$, the product of countably many copies of the space with two points and the discrete topology. By Tikhonov's theorem, Ω is compact. The mapping

$$\phi : \Omega \longrightarrow [0, 1] : (\epsilon_k)_{k \in \mathbb{N}} \mapsto \sum_{k=1}^{\infty} 2^{-k} \epsilon_k$$

(where each ϵ_k is 0 or 1) is a surjection, because each number in $[0, 1]$ has a binary expansion, and is quite easily seen to be continuous *****. By 8.8, this proves that $[0, 1]$ is compact. □

This proof apparently makes essential use of the Axiom of Choice to prove a result that does not need it. However, the particular case of Tikhonov's theorem (countable products of doubletons) that was appealed to can easily be proved without the Axiom. As usual, the Axiom is required to establish the abstract general result for inconceivably large objects, but the "small practical cases" are true anyway. Putting it slightly differently and rather inaccurately, the Axiom is needed when there is no hope of constructing either the desired object (a subcovering, in Tikhonov's theorem) or an explicit contradiction to the denial of its existence.

Lemma 8.25. *A compact set A in $\mathbb{R}$ is bounded.*

Proof. $\{(n-1, n+1) : n \in \mathbb{Z}\}$ is an open covering of $\mathbb{R}$, and so of A; there is a finite subcovering of A, $\{(n_i - 1, n_i + 1) : 1 \leq i \leq k\}$. Then

$$A \subseteq (\min(n_1, \ldots, n_k) - 1, \max(n_1, \ldots, n_k) + 1) + 1),$$

which evidently shows that A is bounded. $\square$

Theorem 8.26. *A set in $\mathbb{R}^n$ is compact if and only if it is closed and bounded.*

At the cost of reiterating the oft-repeated, let me observe that this is a result about $\mathbb{R}^n$ and the standard (Euclidean) metric thereon. It would not be true in most metric spaces, or even in $\mathbb{R}^n$ itself if "boundedness" were understood in terms of some other metric.

Proof. Suppose first that $A \subseteq \mathbb{R}^n$ is compact. Each of the coordinate projections

$$\pi_k : (\xi_1, \xi_2, \ldots, \xi_n) \mapsto \xi_k$$

is continuous, so that $\pi_k(A)$ is compact in $\mathbb{R}$, and bounded: for some $\alpha_k \geq 0$,

$$\pi_k(A) \subseteq [-\alpha_k, \alpha_k].$$

Hence $A \subseteq [-\alpha_1, \alpha_1] \times \cdots \times [-\alpha_n, \alpha_n]$, and is bounded in $\mathbb{R}^n$. Indeed, it is included in the closed ball of radius $\sqrt{\alpha_1^2 + \cdots + \alpha_n^2}$ about the origin.

That A is closed follows instantly from 8.4.

Conversely, if A is closed and bounded, it is a closed subset of the product $[-\alpha, \alpha] \times \cdots \times [-\alpha, \alpha]$ for some suitable $\alpha \geq 0$ (for instance α might be the radius of a ball about the origin including A). The product is compact by Tikhonov's theorem. Apply 8.3. $\square$

Lemma 8.27. *A compact nonnull subset of $\mathbb{R}$ contains its supremum and infimum.* $\square$

Corollary 8.28. *Let $f : \Omega \longrightarrow \mathbb{R}$ be a continuous real-valued function on the topological space Ω. If A is a compact subset of Ω, then*

$$(\exists a, b \in A)(\forall x \in A) \quad f(a) \leq f(x) \leq f(b). \qquad \square$$

This is normally stated as "a continuous function on a compact set is bounded and attains its bounds", "bound" being an abbreviation for "least upper bound" or "greatest lower bound" as the case may be. By "attaining" a "bound" is meant that the "bound" is itself a value of the function. In general this will not be true; for instance, the function $f(t) := t$ is bounded on the open interval $(0, 1)$, with infimum 0 and supremum 1, neither of which is a value of f on the interval.

The attaining of bounds is obviously an important question, not least because some differential equations of physics and geometry have solutions that are critical points of related scalar-valued functions. In such cases the compactness of the domain may be quite difficult to establish. But it is also possible sometimes to "put the compactness into the function" in some way.

9. Compactifications.

How special are compact spaces? The very special properties they have, such as 8.10, might suggest they are very rare. This is not quite the case. Let us begin simply.

Definition 9.1. A topological space Ω is locally compact if every point of Ω has a base of neighbourhoods consisting of compact sets. [Naturally enough, one often speaks of a base of compact neighbourhoods.]

As with the definitions of "regular" and "normal", Kelley had a slightly different definition, requiring each point to have only one compact neighbourhood; and, as before, this is not silly, but has not been generally accepted.

Lemma 9.2. *A locally compact Hausdorff space is regular.* □

Kelley's version was that every point in a regular "Kelley-locally-compact" space has a base of compact neighbourhoods. *****

Suppose that Ω is any set. Then the axioms of set theory imply that there is some element $*$ that does not belong to Ω (we know, for instance, by Cantor's theorem, that $\mathcal{P}(\Omega)$ cannot be a subset of Ω). It is not very important what we take for $*$, provided that it is a "new" element. But let $\Omega^* := \Omega \cup \{*\}$, Ω with one extra element.

Theorem 9.3. *Let Ω be a locally compact Hausdorff space. There is a unique compact Hausdorff topology on Ω^* such that Ω is an open subset and the original topology on Ω is the subspace topology.*

Proof. Define the topology on Ω^* by requiring the "open" sets to be either open subsets of Ω itself (these are the "open" sets not containing $*$) or complements in Ω^* of compact subsets of Ω (these are the "open" sets containing $*$, and meet Ω in open sets of Ω).

$\emptyset$ is an open set of Ω, Ω^* is the complement in Ω^* of the compact set $\emptyset$ in Ω. If U_1, U_2 are open sets of Ω and K_1, K_2 are compact sets of Ω, then K_1, K_2 are closed in Ω by 8.4, so $U_1 \cap (\Omega^* \setminus K_1) = U_1 \setminus K_1$ is open in Ω; $U_1 \cap U_2$ is open in Ω; and $(\Omega^* \setminus K_1) \cap (\Omega^* \setminus K_2) = \Omega^* \setminus (K_1 \cup K_2)$ is again the complement in Ω^* of a compact set in Ω, by 8.6. Finally, for any class $\mathcal{U}$ of "open" sets in Ω^*, there are two possibilities. Either they are all open subsets of Ω; then their union is also an open subset of Ω, and so an "open" set of Ω^*; or else at least one of them contains $*$ and has a complement compact in Ω, in which case $* \in \bigcup_{U \in \mathcal{U}} U$ and

$$\Omega^* \setminus \left(\bigcup_{U\in\mathcal{U}} U\right) = \bigcap_{U\in\mathcal{U}} (\Omega^* \setminus U)$$

is the intersection of a compact subset of Ω with closed subsets of Ω, so is compact by 8.3. In either case, the union is an "open" set of Ω^*. This completes the proof that the "open" sets form a topology on Ω^*. Evidently Ω is itself "open", and its original topology is the same as the subspace topology.

Any "open" covering $\mathcal{V}$ of Ω^* must have one member V containing $*$. As $\Omega^* \setminus V$ is a compact subset of Ω, and is covered by $\{U\cap\Omega : U\in\mathcal{V}\}$, there are finitely many members $U_1, U_2, \ldots, U_n$ of $\mathcal{V}$ such that $\Omega^* \setminus V \subseteq \bigcup_{i=1}^n (U_i \cap \Omega)$; hence

$$\Omega^* \subseteq V \cup \left(\bigcup_{i=1}^n U_i\right),$$

so that $\mathcal{V}$ has a finite subcovering. Hence, Ω^* is compact.

Ω^* is Hausdorff. If $x, y \in \Omega$ and $x \neq y$, then there are open sets of Ω, U and V, such that $x \in U$, $y \in V$, and $U \cap V = \emptyset$. But U and V are also open sets in the topology of Ω^*. If $x \in \Omega$, there is, by hypothesis, a compact neighbourhood K of x in Ω: $x \in \mathrm{int}_\Omega K$, which is open in Ω^*, and $* \in \Omega^* \setminus K$, which is also open in Ω^*. These "open" sets are disjoint. Thus, Ω^* is Hausdorff.

That the topology specified is unique subject to the stated conditions is easy. □

Definition 9.4. The space Ω^* just constructed is called the Aleksandrov compactification or one-point compactification of the Hausdorff locally compact space Ω.

Definition 9.5. Let Ω be any topological space. A compactification of Ω is a mapping $k : \Omega \longrightarrow X$, where X is a second topological space, such that

(a) X is compact, (b) $k(\Omega)$ is dense in X, i.e. $\mathrm{cl}_X(k(\Omega)) = X$, and

(c) k is a homeomorphism of Ω with $k(\Omega)$, for the subspace topology on $k(\Omega)$ induced from X.

I shall consider seriously only compactifications for which the compactified space X is T_2, because of 8.9 (see below), but the definition can be given more generally. The condition (b) ensures that X is not bigger than it need be. The Aleksandrov compactification has the property that the "embedding" k is an open continuous mapping, which is not usually the case.

Example 9.6. Let $\Omega := \mathbb{R}$, which is Hausdorff locally compact. The one-point compactification of Ω, which I shall call Ω_1, is homeomorphic with the circle $S^1 := \{(x,y) \in \mathbb{R}^2 : x^2 + y^2 = 1\}$. The embedding $k_1 : \Omega \longrightarrow S^1$ may be

$$t \mapsto \left(\frac{1-t^2}{1+t^2}, \frac{2t}{1+t^2}\right) = (\cos(2\tan^{-1}t), \sin(2\tan^{-1}t)).$$

The "ideal point" or "point at infinity" $*$ corresponds to $(-1, 0)$.

There is another well-known Hausdorff compactification of Ω, a "two-point compactification" Ω_2 homeomorphic with $[-1, 1]$. $k_2 : \Omega \longrightarrow [-1, 1]$ may be taken as $t \mapsto 2(\tan^{-1}t)/\pi$. Thus $k_2(\Omega)$ is the open interval $(-1, 1)$.

There is a continuous mapping $\varpi : \Omega_2 \longrightarrow \Omega_1$ defined by

$$\varpi(s) := (\cos(\pi t), \sin(\pi t)) ,$$

and $\varpi \circ k_2 = k_1$.

There are other possible non-Hausdorff compactifications of $\mathbb{R}$. One might, for instance, define $X := \Omega \cup \{*, \dagger\}$, with the open sets being defined as

the open sets of Ω; and
the sets of X that contain $*$, or $\dagger$, or both,
but meet Ω in the complement of a compact set.

Then X is compact, but not T_1.

The existence of a T_2 two-point compactification of $\mathbb{R}$ is a consequence of the fact that $\mathbb{R}$ has two "ends" (a technical term).

Example 9.7. Let $\Omega := \mathbb{R}^n$. Then the Aleksandrov compactification of Ω is homeomorphic with S^n. ****

On the other hand, for $n > 1$, Ω is homeomorphic to the unit ball $B(0; 1)$ in $\mathbb{R}^n$ with respect to the Euclidean distance. But the closed ball $C(0; 1)$ is compact; we may regard it as a Hausdorff compactification of $B(0; 1)$, with the inclusion as the embedding "k". In this way we obtain a compactification of Ω in which the ideal points form an $(n-1)$-dimensional sphere.

Projective n-space is another Hausdorff compactification of Ω. In this case, the ideal points, the "points at infinity", form a copy of projective $(n-1)$-space. ****

9.5 has one consequence that is worth noting. The *construction* of the one-point compactification may still be performed if the locally compact Hausdorff space Ω is in fact compact to begin with; the ideal point $*$ is then both open and closed, and the inclusion of Ω in Ω^* is not a compactification, because 9.5*(b)* does not hold. More generally, a compact Hausdorff space may be characterized as a space for which the only Hausdorff compactifications are homeomorphisms (i.e. for which k is onto).

These various examples suggest a partial order.

Definition 9.8. Let Ω be a topological space, and suppose that $k_1 : \Omega \longrightarrow X_1$ and $k_2 : \Omega \longrightarrow X_2$ are Hausdorff compactifications of Ω. k_1 dominates k_2 , and we write $k_1 \geq k_2$, if there is a continuous mapping $\varpi_{12} : X_1 \longrightarrow X_2$ such that $\varpi_{12} \circ k_1 = k_2$.

Since I did not introduce these ideas earlier, let me have here a tiny digression.

Definition 9.9. Let A and B be subsets of the topological space Ω. A is dense in B if $B \subseteq \mathrm{cl}_\Omega(A)$. [Commonly $A \subseteq B$ when this nomenclature is used.]

Lemma 9.10. *Suppose Ω and Ψ are topological spaces, Ψ is Hausdorff, A and B are subsets of Ω, and $f, g : \Omega \longrightarrow \Psi$ are continuous. If A and B are subsets of Ω and A is dense in B and $f|A = g|A$, then $f|B = g|B$.*

In words: continuous extension from dense subsets into Hausdorff spaces is unique.

Proof. Suppose $x \in B$ and $f(x) \neq g(x)$. Then there are open sets U, V in Ψ such that $f(x) \in U$, $g(x) \in V$, and $U \cap V = \emptyset$. But $f^{-1}(U) \cap g^{-1}(V)$ is an open neighbourhood of x in Ω; as such, it contains a point y of A; and, by hypothesis, $f(y) = g(y)$, which is absurd, since $f(y) \in U$ and $g(y) \in V$. □

Of course, not every continuous mapping defined on A with its subspace topology will extend to a continuous mapping defined on Ω.

Lemma 9.11. *If a mapping ϖ_{12} exists as in* 9.8, *it is unique. It must be a surjection, and the topology on X_2 must be the topology induced from X_1 via ϖ_{12}. If $k_2 \geq k_1$ as well, then ϖ_{12} is a homeomorphism with inverse ϖ_{21}.*

Proof. The image of ϖ_{12} must be compact, and so closed, in X_2, and includes $k_2(\Omega)$. Thus by 9.5*(b)* it must be the whole of X_2. The assertion about the induced (quotient) topology follows from 8.10. As $\varpi_{12} \circ k_1 = k_2$, ϖ_{12} is completely determined on $k_1(\Omega)$, and, as this is dense in X_1, the uniqueness of ϖ_{12} follows from 9.10. If $\varpi_{21} : X_2 \longrightarrow X_1$ also exists (with $\varpi_{21} \circ k_2 = k_1$), then $\varpi_{21} \circ \varpi_{12} \circ k_1 = \varpi_{21} \circ k_2 = k_1$, and uniqueness shows that $\varpi_{21} \circ \varpi_{12}$ is the identity of X_1; similarly, $\varpi_{12} \circ \varpi_{21}$ is the identity of X_2. Thus, ϖ_{21} and ϖ_{12} are inverse homeomorphisms. □

In effect, if $k_1 \geq k_2$ and $k_2 \geq k_1$, the two Hausdorff compactifications are "equivalent". For instance, a noncompact Hausdorff locally compact space has essentially only one one-point compactification — which, of course, is obvious anyway.

Remark 9.12. It is clear that the relation $\geq$ between Hausdorff compactifications is transitive and reflexive; it is a "weak" partial order (not antisymmetric).

Lemma 9.13. *Suppose that $k_\alpha : \Omega \longrightarrow X_\alpha$ are Hausdorff compactifications, for α in an index set A. Then there is a Hausdorff compactification $k : \Omega \longrightarrow X$ such that $k \geq k_\alpha$ for each $\alpha \in A$.*

Proof. Let $Y := \prod_{\alpha \in A} X_\alpha$, which is compact T_2. Define

$$k : \Omega \longrightarrow Y : x \mapsto (k_\alpha(x))_{\alpha \in A}$$

and $\pi_\alpha : Y \longrightarrow X_\alpha$ is the coordinate projection, as usual. Clearly $\pi_\alpha \circ k = k_\alpha$ for each α, and it follows that k is continuous and injective. Let $X := \mathrm{cl}_Y(k(X))$. Then X is compact T_2, k maps into X, and $k(\Omega)$ is dense in X.

Let U be open in Ω. Then (taking some $\alpha \in A$) $k_\alpha(U)$ is, by hypothesis, relatively open in $k_\alpha(\Omega)$, so there is some V open in X_α such that $k_\alpha(\Omega) \cap V = k_\alpha(U)$. But then $k(\Omega) \cap \pi_\alpha^{-1}(V) = k(U)$, which shows that $k(U)$ is relatively open in $k(\Omega)$. This completes the proof. □

Of course we do not yet know that a Hausdorff compactification of Ω exists. However, it might appear from 9.13 that, if any Hausdorff compactification exists, there must be a greatest one. Unfortunately, there is also the set-theoretic problem that the class of all Hausdorff compactifications is not obviously a set — and, in principle, the construction may therefore not give a set either. There is, however, a construction which not only avoids this difficulty, but also has an interesting extension property.

Remark 9.14. If Ω has a Hausdorff compactification, $k : \Omega \longrightarrow X$, then X is normal, and so its subspace $k(\Omega)$ is completely regular (Tikhonov). But, as k is a homeomorphism of Ω with its image, Ω itself must be Tikhonov.

Conversely, suppose Ω is completely regular. Let $F(\Omega)$ denote the class of all continuous functions $f : \Omega \longrightarrow I$, and set $\mathfrak{X}_\Omega := \prod_{f \in F(\Omega)} I$. By Tikhonov's theorem, $\mathfrak{X}_\Omega$ is compact Hausdorff, and there is a mapping $k_\Omega : \Omega \longrightarrow \mathfrak{X}_\Omega$:

$$(\forall x \in \Omega) \quad k_\Omega(x) := (f(x))_{f \in F(\Omega)} .$$

k is continuous (as $\pi_f \circ k_\Omega = f$ is continuous for each f). It is one-one, since, if $x \neq y$ in Ω, there is some $f \in F(\Omega)$ with $f(x) \neq f(y)$.

I claim that k_Ω is a homeomorphism of Ω with $k_\Omega(\Omega)$. Given an open set U of Ω and any $x \in U$, there is some $f \in F(\Omega)$ such that $V_f := f^{-1}(-\infty, 1) \subseteq U$. Then $k_\Omega(V_f) = k_\Omega(\Omega) \cap \pi_f^{-1}(V_f)$, and so is relatively open in $k_\Omega(\Omega)$. As this holds for each $x \in U$, $k_\Omega(U)$ is relatively open in $k_\Omega(\Omega)$. Hence, k_Ω is a homeomorphism with $k_\Omega(\Omega)$.

Finally, define $X := \mathrm{cl}_{\mathfrak{X}_\Omega}(k_\Omega(\Omega))$. Then $k_\Omega : \Omega \longrightarrow X$ is a Hausdorff compactification of Ω. It is called the Stone-Čech compactification, customarily denoted $\beta\Omega$.

Lemma 9.15. *If Ω is compact, then $k_\Omega : \Omega \longrightarrow \beta\Omega$ is a homeomorphism.*

Proof. In this case $k_\Omega(\Omega)$ is compact, so it is closed in $\mathfrak{X}_\Omega$ and $k_\Omega(\Omega) = \beta\Omega$. ☐

Lemma 9.16. *If $\phi : \Omega \longrightarrow \Psi$ is a continuous mapping between Tikhonov spaces, there is a unique continuous mapping $\psi : \beta\Omega \longrightarrow \beta\Psi$ such that the diagram*

$$\begin{array}{ccc} \Omega & \xrightarrow{\phi} & \Psi \\ k_\Omega \downarrow & & \downarrow k_\Psi \\ \beta\Omega & \xrightarrow[\psi]{} & \beta\Psi \end{array} \tag{24}$$

commutes. [It follows easily that β is a covariant functor from the category of Tikhonov spaces and continuous maps to the category of compact Hausdorff spaces and continuous maps, and the embeddings k_Ω constitute a natural transformation between the identity functor and β. It would be natural to write $\beta\phi$ instead of ψ.]

Proof. There is a mapping $F(\phi) : F(\Psi) \longrightarrow F(\Omega) : g \mapsto g \circ \phi$. (I am not proposing any topology on $F(\Psi)$ or $F(\Omega)$; this is just a set-theoretic construction as far as we are concerned. Notice that it is *contravariant,* i.e. goes in the opposite direction from ϕ.) In turn, there is an induced mapping $\Phi : \mathfrak{X}_\Omega \longrightarrow \mathfrak{X}_\Psi$,

$$\Phi((t_f)_{f \in F(\Omega)}) := (t_{F(\phi)g})_{g \in F(\Psi)} .$$

This formula is more comprehensible if we think of $\mathfrak{X}_\Omega, \mathfrak{X}_\Psi$ as sets of "choice functions" (§4**A**). If $s : F(\Omega) \longrightarrow I$ is a choice function, then $\Phi(s) = s \circ F(\phi)$. Φ is a continuous mapping, since its composite with any coordinate projection $\pi_g \circ \Phi$ of $\mathfrak{X}_\Psi$ is the coordinate projection $\pi_{g \circ \phi}$ of $\mathfrak{X}_\Omega$. Notice that Φ is *covariant*, i.e. goes in the same direction as ϕ. It is easily checked that

$$\begin{array}{ccc} \Omega & \xrightarrow{\phi} & \Psi \\ k_\Omega \downarrow & & \downarrow k_\Psi \\ \mathfrak{X}_\Omega & \xrightarrow[\Phi]{} & \mathfrak{X}_\Psi \end{array} \tag{25}$$

commutes, by the calculation:

$$\begin{aligned} \Phi \circ k_\Omega(x) &= \Phi((f(x))_{f\in F(\Omega)}) = ((F(\phi)g)(x))_{g\in F(\Psi)} \\ &= (g(\phi(x)))_{g\in F(\Psi)} = k_\Psi(\phi(x)) \,. \end{aligned}$$

As Φ is continuous, $\Phi(\mathrm{cl}_{\mathfrak{X}_\Omega}(k_\Omega(\Omega))) \subseteq \mathrm{cl}_{\mathfrak{X}_\Psi}(k_\Psi(\Psi))$, so that Φ restricts to a continuous mapping $\psi : \beta\Omega \longrightarrow \beta\Psi$. The commutative diagram (24) results from (25). Uniqueness is assured by the fact that $k_\Omega(\Omega)$ is dense in $\beta\Omega$, 9.10. □

Proposition 9.17. *Let Ω be a Tikhonov space, and let $k_\Omega : \Omega \longrightarrow \beta\Omega$ be its Stone-Čech compactification. If Ψ is any compact Hausdorff space and $\phi : \Omega \longrightarrow \Psi$ is continuous, there is a unique continuous mapping $\psi : \beta\Omega \longrightarrow \Psi$ such that $\psi \circ k_\Omega = \phi$.*

Proof. In this case $k_\Psi : \Psi \longrightarrow \beta\Psi$ is a homeomorphism, by 9.15. So "ψ" here may be what in the notation of 9.16 would be called $k_\Psi^{-1} \circ \psi$. □

Theorem 9.18. *The Stone-Čech compactification $k_\Omega : \Omega \longrightarrow \beta\Omega$ of a completely regular space Ω has the following properties.*

(a) *It dominates any other Hausdorff compactification of Ω.*

(b) *Any continuous mapping of Ω into a compact Hausdorff space Ψ has a unique continuous extension to a continuous mapping of $\beta\Omega$ into Ψ.*

(c) *Any bounded real-valued continuous function on Ω has a unique extension to a continuous real-valued function on $\beta\Omega$ (having the same supremum and infimum).*

(d) *Any Hausdorff compactification having either of the properties* (a) *or* (b) *is equivalent to the Stone-Čech compactification.*

Proof. *(b)* is 9.17. If $k : \Omega \longrightarrow X$ is any Hausdorff compactification, then, by *(b)*, there is a unique $\varpi : \beta\Omega \longrightarrow X$ such that $k = \varpi \circ k_\Omega$; this proves *(a)*. Given a bounded continuous function $f : \Omega \longrightarrow \mathbb{R}$, let $\Psi := [\inf f(\Omega), \sup f(\Omega)]$, in the subspace topology; then *(c)* follows from *(b)*. *(d)* is an obvious consequence of 9.11. □

The Stone-Čech compactification clearly has a privileged position among Hausdorff compactifications of Tikhonov spaces, but for more restricted classes of spaces and for specific problems there are other possible compactifications. (A notable example is Carathéodory's theory of the boundary behaviour of conformal transformations of regions in $\mathbb{C}$. I pointed out at 9.7 that the closed unit disk in $\mathbb{C}$ is a compactification of the open unit disk; there are many others.

Although $\beta\Omega$ is obviously important, it is rather obscure. Even in simple cases, it lacks a simple and attractive description. For instance the space $\beta\mathbb{N}$ occasionally appears, and is described simply as $\beta\mathbb{N}$.

10. Connectedness.

Definition 10.1. A topological space Ω is *connected* if $\emptyset$ and Ω are the only subsets of Ω that are both open and closed. A subset M of Ω is connected if it is connected in the subspace topology.

Lemma 10.2. *M is connected if, whenever U, V are open sets of Ω such that $M \subseteq U \cup V$ and $U \cap V \cap M = \emptyset$, either $M \subseteq U$ or $M \subseteq V$; or, equivalently, if, whenever C, D are closed sets of Ω such that $M \subseteq C \cup D$ and $C \cap D \cap M = \emptyset$, either $M \subseteq C$ or $M \subseteq D$.* □

Recall from 7.10 that subsets A, B of Ω are "separated" if

$$A \cap \mathrm{cl}_\Omega(B) = B \cap \mathrm{cl}_\Omega(A) = \emptyset .$$

Lemma 10.3. *M is connected if and only if it cannot be expressed as the union of two non-empty separated sets.* [The proof of 7.10 is relevant.] □

Lemma 10.4. *Let $f : \Omega \longrightarrow \Psi$ be a continuous mapping between topological spaces. If M is a connected subset of Ω, $f(M)$ is a connected subset of Ψ.* □

In this respect, connectedness is like compactness. As with compactness, I have not given any example so far.

Proposition 10.5. *Any closed bounded interval in $\mathbb{R}$ is connected.*

Proof. As before, it will suffice to consider $I = [0, 1]$. Suppose there are open sets U, V in $\mathbb{R}$ such that $U \cap V \cap I = \emptyset$ and $I \subseteq U \cup V$, and assume that $0 \in U$. Let

$$S := \{t \in I : [0, t] \subseteq U\} \neq \emptyset , \quad s := \sup S .$$

(S is of course bounded above by 1.)

I claim that $s \in U \cap I$. This is obvious, if we note that $U \cap I$ is relatively closed in I, and so closed in $\mathbb{R}$, and that s, as the supremum of $S \subseteq U \cap I$, is an adherent point of S and so of $U \cap I$. But here is an (equivalent) elementary proof.

Certainly $0 \leq s \leq 1$. If $s \notin U \cap I$, then $0 < s \in V \cap I$. But V is open in $\mathbb{R}$; so there is some $\delta > 0$ such that $(s - \delta, s + \delta) \subseteq V$, and $(s - \delta, s + \delta) \cap I \subseteq V \cap I$. This is absurd, since it implies that $S \subseteq [0, \max(0, s - \delta)]$. Thus, indeed, $s \in U \cap I$.

However, there is some $\epsilon > 0$ such that $(s - \epsilon, s + \epsilon) \subseteq U$; if $s < 1$, it follows that $[s, s + \min(\frac{1}{2}\epsilon, 1 - s)] \subseteq U \cap I$. Then $[s, s + \min(\frac{1}{2}\epsilon, 1 - s)] \subseteq S$, and this contradicts the definition of s as the supremum of s, which is absurd. The conclusion must be that $s = 1$, and, consequently, that $[0, 1] \subseteq U$.

What has just been proved is that, if $0 \in U$, then $I \subseteq U$. But the same proof (*mutatis mutandis*) shows that, if $0 \in V$ instead, then $I \subseteq V$. ☐

Proposition 10.6. *The closure of a connected set is connected.*

Proof. Suppose that M is connected, and that C and D are closed sets such that $\mathrm{cl}_\Omega(M) \subseteq C \cup D$, $C \cap D \cap \mathrm{cl}_\Omega(M) = \emptyset$. Then $M \subseteq C \cup D$, $C \cap D \cap M = \emptyset$. But then either $M \subseteq C$, when, as C is closed, $\mathrm{cl}_\Omega(M) \subseteq C$ too, or $M \subseteq D$, when $\mathrm{cl}_\Omega(M) \subseteq D$. ☐

Lemma 10.7. *Let Ω be a Tikhonov space. Then Ω is connected if and only if $\beta\Omega$ is connected.*

Proof. Let $f : \Omega \longrightarrow \beta\Omega$ be the canonical embedding. If Ω is connected, then so is $f(\Omega)$, and so is $\beta\Omega = \mathrm{cl}_{\beta\Omega}(f(\Omega))$.

Conversely, suppose that Ω is not connected ("disconnected"). Then $\Omega = C \cup D$, where C and D are closed, disjoint, and nonnull. Thus the function $\phi : \Omega \longrightarrow I$ defined by $\phi(C) = \{0\}$, $\phi(D) = \{1\}$, is continuous. By 9.18*(c)*, ϕ extends to a unique continuous function $\widehat{\phi} : \beta\Omega \longrightarrow I$ such that $\widehat{\phi} \circ f = \phi$. But, as $f(\Omega)$ is dense in $\beta\Omega$, 3.14 shows that $\phi(\Omega) \subseteq \widehat{\phi}(\beta\Omega) \subseteq \mathrm{cl}_I(\phi(\Omega))$, and $\widehat{\phi}$ also takes only the values 0 and 1 (and both). Hence $(\widehat{\phi})^{-1}\{0\}, (\widehat{\phi})^{-1}\{1\}$ are closed disjoint nonnull sets whose union is $\beta\Omega$, so that $\beta\Omega$ is also disconnected. ☐

Of course it also follows that $(\widehat{\phi})^{-1}\{0\} = \mathrm{cl}_{\beta\Omega}(f(C))$, and likewise for D. The closure in $\beta\Omega$ of a clopen subset of Ω is clopen in $\beta\Omega$.

Lemma 10.8. *Let $\mathcal{A}$ be any family of nonnull connected subsets of Ω. If no two members of $\mathcal{A}$ are separated, then $M := \bigcup_{A \in \mathcal{A}} A$ is also connected.*

Proof. Suppose that C, D are closed sets in Ω such that $M \subseteq C \cup D$ and $C \cap D \cap M = \emptyset$. For any $A \in \mathcal{A}$, $A \subseteq C \cup D$ and $C \cap D \cap A = \emptyset$, and so either $A \subseteq C$ or $A \subseteq D$, but not both (as $A \neq \emptyset$). Thus $\mathcal{A}$ is the disjoint union of two subclasses, $\mathcal{A}_1 := \{A \in \mathcal{A} : A \subseteq C\}$ and $\mathcal{A}_2 := \{A \in \mathcal{A} : A \subseteq D\}$.

However, let $A_1 \in \mathcal{A}_1$ and $A_2 \in \mathcal{A}_2$. Then $\mathrm{cl}(A_1) \subseteq C$, and consequently

$$\mathrm{cl}(A_1) \cap A_2 \subseteq C \cap (D \cap M) = \emptyset .$$

Similarly, $A_1 \cap \mathrm{cl}(A_2) = \emptyset$. This is contrary to the hypothesis that A_1, A_2 are not separated. It must, therefore, be the case that either $\mathcal{A}_1$ or $\mathcal{A}_2$ is empty, and the other is $\mathcal{A}$. If we suppose that $\mathcal{A}_2 = \emptyset$, then $M = \bigcup_{A \in \mathcal{A}_1} A \subseteq C$. If $\mathcal{A}_1 = \emptyset$, then $M \subseteq D$. Thus, the condition of 10.2 is satisfied. ☐

In particular, if $\bigcap_{A \in \mathcal{A}} A \neq \emptyset$, then $\bigcup_{A \in \mathcal{A}} A$ is connected.

Definition 10.9. Say [unofficially] that two *points* $x, y \in \Omega$ are "connected" in Ω if there is some connected subset M of Ω such that $x, y \in M$. This is clearly a reflexive and symmetric relation in Ω; but 10.8 shows that it is also transitive. Thus it is an equivalence relation $\sim$ (we may [unofficially] call it the "connectedness relation").

The equivalence classes of Ω under the connectedness relation $\sim$ are called the *components* (or, if there is any ambiguity — and there quite often is — the *connected components*) of Ω. [The components of a subset of Ω are its components with respect to the subspace topology.] The component in Ω of a point a is its $\sim$-equivalence class. This is not, perhaps, the standard definition, which is stated below:

Lemma 10.10. *The component in Ω of $a \in \Omega$ is the largest connected subset of Ω that contains a. Indeed, it is the union of all the connected subsets that contain a.* □

Lemma 10.11. *Each component of Ω is closed.*

Proof. Since the closure of a component is also connected by 10.6, it must be equal to the component itself, by 10.10. □

Proposition 10.12. *Suppose that the topological space Ω_β is connected for each index β in the index set B. Then $\Omega := \prod_{\beta \in B} \Omega_\beta$ is connected.*

Proof. Suppose that $(u_\beta), (v_\beta)$ are two points of the product. If they differ in only the γth coordinate ($u_\beta = v_\beta$ unless $\beta = \gamma$), they are connected in Ω, since the subset $\{(x_\beta) \in \Omega : \beta \neq \gamma \Rightarrow x_\beta = u_\beta\}$ of Ω is homeomorphic to Ω_γ and so connected. But, by transitivity, it follows that the component of the point $u := (u_\beta)$ includes the set of all points differing from u in only finitely many coordinates. This set, however, is dense in Ω (because of the definition of the product topology). By 10.11, the component of u must be the whole of Ω. □

If M_β is a subset of Ω_β for each β, then $M := \prod_{\beta \in B} M_\beta$ is a subset of $\Omega := \prod_{\beta \in B} \Omega_\beta$, and the subspace topology on M is the same as the product of the subspace topologies on the M_β. This remark offers a version of 10.12 for subsets.

If Ω is expressed as the disjoint union of two open (and closed) sets U and V, and Q is a component of Ω, then necessarily $Q \subseteq U$ or $Q \subseteq V$. It is tempting to suppose that the "converse" of this statement holds: that the component of a point $a \in \Omega$ is the intersection of all the open-and-closed sets containing a. This is not true in general***. In fact its falsity is expressed in a definition:

Definition 10.13. Let $a \in \Omega$. The *quasicomponent* of a in Ω is the intersection of all the open-and-closed sets containing a.

We can define two points $a, b \in \Omega$ to be *quasiconnected* if every clopen set containing a contains b, and *vice versa*. This is an equivalence relation on Ω.***

Lemma 10.14. (a) *The quasicomponents of the points of Ω form a partition of Ω by closed sets.*

(b) *For each $a \in \Omega$, the component of a in Ω is included in the quasicomponent of a in Ω.*

(c) *Each quasicomponent is a union of some components of Ω.* □

Example 10.15. Let A be the subset of $\mathbb{R}^2$ consisting of the points $(0,0)$ and $(0,1)$ and the line segments $\{(\frac{1}{n},y) : 0 \leq y \leq 1\}$. Evidently the components of A are the singletons $\{(0,0)\}$ and $\{(0,1)\}$ and the individual line segments. The line segments are also quasicomponents, but the quasicomponent of $(0,0)$ is the doubleton $\{(0,0),(0,1)\}$. This is because any (relatively) open set that contains $(0,0)$ must also meet all the segments for sufficiently large n, but if it is also (relatively) closed it must then include all these segments, and in consequence must also contain $(0,1)$.

Theorem 10.16. *Let Ω be a compact Hausdorff space, and $a \in \Omega$. Then the component of a in Ω is the same as its quasicomponent.*

Proof. All we have to show is that the quasicomponent Q of a is connected. Q is closed, so it will suffice to suppose there are nonnull closed sets C, D in Ω such that $Q = C \cup D$ and $C \cap D = \emptyset$. We may suppose $a \in C$. Ω is normal by 8.5, so there are open sets U, V such that $C \subseteq U$, $D \subseteq V$, and $U \cap V = \emptyset$, and $Q \subseteq U \cup V$.

By definition, $Q = \bigcap_{X \in \mathfrak{X}} X$, where $\mathfrak{X}$ is the class of open-and-closed sets containing a. The intersection of any finite number of members of $\mathfrak{X}$ is still in $\mathfrak{X}$.

Suppose, if possible, that $X \setminus (U \cup V) \neq \emptyset$ for each $X \in \mathfrak{X}$. These sets are all closed, and, by the hypothesis, they form a system with the finite intersection property (see 8.7). As Ω is compact, $\bigcap_{X \in \mathfrak{X}} (X \setminus (U \cup V)) \neq \emptyset$ too, by 8.7. This says that $Q \setminus (U \cup V) \neq \emptyset$, which is absurd. The conclusion must be that there is some $X \in \mathfrak{X}$ for which $X \subseteq U \cup V$.

In that case, however, $X \cap U = X \setminus V$ is also open-and-closed, and it contains a (recall that $a \in C \subseteq U$ and $a \in Q \subseteq X$). So $Q \subseteq X \cap U$; which implies that $Q \cap D \subseteq Q \cap V = \emptyset$, and, consequently, that $Q \subseteq C$. This shows that Q is indeed connected. $\square$

Although there is quite a lot more that could be said about connectedness, I shall finish this section with a few definitions of concepts that are often enough mentioned to deserve elucidation, even though there is no reason to discuss them further here.

Definition 10.17. A topological space Ω is described as a *continuum* if it is compact and connected.

Definition 10.18. Ω is *locally connected* if each point of Ω has a base of connected neighbourhoods.

Definition 10.19. Ω is *totally disconnected* if all its connected components are singletons. [Some people call this *hereditarily disconnected*, a rather silly name.]

Definition 10.20. Ω is *zero-dimensional* if its topology has a base of open-and-closed sets.

Definition 10.21. Ω is *extremally disconnected* if the closure of every open set of Ω is open. [I have seen the phrase *extremely disconnected*, but suspect it is a typist's error.]

11 Paracompactness.

Paracompactness is where the geometers and the analysts part company. Let me start with the historical background.

Remark 11.1. Let (Ω, d) be a separable metric space, with a countable dense subset Y. The class $\{B(y; \frac{1}{n}) : y \in Y \;\;\&\;\; n \in \mathbb{N}\}$ of open balls is a base for the topology of Ω; indeed, if U is an open set in Ω and $x \in U$, there is some $\delta > 0$ such that $B(x; \delta) \subseteq U$, and some $n \in \mathbb{N}$ such that $\frac{1}{n} \leq \delta$. Take $y \in B(x; \frac{1}{2n}) \cap Y$, and then

$$x \in B(y; \tfrac{1}{2n}) \subseteq B(x; \tfrac{1}{n}) \subseteq U\,.$$

(Both the symmetry of the metric and the triangle inequality have been used.) Thus, in fact, Ω is second countable, with the countable base consisting of open balls.

Let $\mathcal{U}$ be any open covering of Ω. Then $\mathcal{U}$ has a countable *refinement* consisting of open balls — in fact, of those balls that belong to the base just constructed and are included in a member of $\mathcal{U}$. [I recall that a covering $\mathcal{V}$ *refines* or *is a refinement of* a covering $\mathcal{U}$ if every member of $\mathcal{V}$ is a subset of some member of $\mathcal{U}$. It is clear that a refinement of a refinement is a refinement.]

Now suppose that $\{B(a_i; r_i) : i \in \mathbb{N}\}$ is any covering of Ω by open balls, where each $r_i > 0$. Define, for each $k \in \mathbb{N}$,

$$W_k := B(a_k; r_k) \setminus \left(\bigcup_{i=1}^{k-1} C(a_i; r_i - \tfrac{1}{k})\right),$$

where $C(a_i; r_i - \frac{1}{k})$ denotes the closed metric ball $\{x \in \Omega : d(x, a_i) \leq r_i - \frac{1}{k}\}$, which is a closed set in Ω (and may be empty, if $r_i < \frac{1}{k}$). Hence, W_k is also an open set in Ω.

More surprisingly, $\{W_k : k \in \mathbb{N}\}$ is also an open covering of Ω. If $x \in \Omega$, certainly there exists some k for which $x \in B(a_k; r_k)$, and, that being so, there must exist a *least* natural number k for which $x \in B(a_k; r_k)$. But then $x \in W_k$.

However, more is true. As $x \in B(a_k; r_k)$, $d(x, a_k) < r_k$, and there exists $l \in \mathbb{N}$ such that $2/l \leq r_k - d(x, a_k)$; this ensures that, whenever $l \leq m \in \mathbb{N}$,

$$B(x; \tfrac{1}{l}) \subseteq B(a_k; r_k - \tfrac{1}{m}) \subseteq C(a_k; r_k - \tfrac{1}{m})\,,$$

by the triangle inequality. Thus $B(x; \frac{1}{l}) \cap W_m = \emptyset$ when $m \geq l$. The effect of this is that, for any point of Ω, there is a neighbourhood of x (the ball $B(x; \frac{1}{l})$) which meets only finitely many of the sets W_k.

The argument just given is due to Dieudonné, who also proposed the most important definitions that follow.

Definition 11.2. Let $\mathcal{C}$ be any class of subsets of the topological space Ω. $\mathcal{C}$ is *locally finite* if, for each $x \in \Omega$, there exists $M \in \mathfrak{N}(x)$ such that $\{C \in \mathcal{C} : M \cap C \neq \emptyset\}$ is finite. $\mathcal{C}$ is *discrete* if, for each $x \in \Omega$, there exists $M \in \mathfrak{N}(x)$ such that $M \cap C = \emptyset$ for all $C \in \mathcal{C}$ except at most one; that is, $\{C \in \mathcal{C} : M \cap C \neq \emptyset\}$ is empty or a singleton. [Notice that here we are speaking of a discrete class of subsets — a discrete subset is something different.]

The class $\mathcal{C}$ is σ-locally finite if it can be expressed as a union $\mathcal{C} = \bigcup_{k=1}^{\infty} \mathcal{C}_n$ in which each summand $\mathcal{C}_n$ is locally finite; it is σ-discrete if it can be expressed as a countable union of discrete classes.

Notice that any subclass of a locally finite [*or* discrete *or* σ-discrete *or* σ-locally finite] class has the same property.

Lemma 11.3. *Let $\mathcal{C}$ be a locally finite class of sets. Then $\{\mathrm{cl}(C) : C \in \mathcal{C}\}$ is also locally finite, and*

$$\mathrm{cl}\Big(\bigcup_{C\in\mathcal{C}} C\Big) = \bigcup_{C\in\mathcal{C}} \mathrm{cl}(C)\,.$$

In particular, the union of a locally finite class of closed sets is also closed.

Proof. Suppose that x is an adherent point of $\bigcup_{C\in\mathcal{C}} C$. As $\mathcal{C}$ is locally finite, there is some $M \in \mathfrak{N}(x)$ such that $\{C \in \mathcal{C} : C \cap M \neq \emptyset\}$ is finite, being $\{C_1, C_2, \ldots, C_k\}$ say. Hence, $M \cap (\bigcup_{C\in\mathcal{C}} C) = M \cap \Big(\bigcup_{i=1}^{k} C_i\Big)$.

Let $N \in \mathfrak{N}(x)$. Then $M \cap N \in \mathfrak{N}(x)$, and by the definition of adherent point

$$M \cap N \cap \Big(\bigcup_{i=1}^{k} C_i\Big) = M \cap N \cap \Big(\bigcup_{C\in\mathcal{C}} C\Big) \neq \emptyset\,.$$

Hence, x is an adherent point of $\bigcup_{i=1}^{k} C_i$. Thus

$$x \in \mathrm{cl}\Big(\bigcup_{i=1}^{k} C_i\Big) = \bigcup_{i=1}^{k} \mathrm{cl}(C_i) \subseteq \bigcup_{C\in\mathcal{C}} \mathrm{cl}(C)\,,$$

and this shows that $\mathrm{cl}(\bigcup_{C\in\mathcal{C}} C) \subseteq \bigcup_{C\in\mathcal{C}} \mathrm{cl}(C)$. The converse inclusion is obvious. □

Given $x \in \Omega$ and $M \in \mathfrak{N}(x)$, let

$$\mathcal{C}(M) := \{C \in \mathcal{C} : M \cap C \neq \emptyset\}\,. \tag{26}$$

Obviously $M_1 \subseteq M_2 \Longrightarrow \mathcal{C}(M_1) \subseteq \mathcal{C}(M_2)$. If there is a neighbourhood M of x such that $\mathcal{C}(M)$ is finite, then there must be a neighbourhood $N \subseteq M$ such that $\mathcal{C}(N)$ has the least possible number of members. Then, for each $C \in \mathcal{C}(N)$, $x \in \mathrm{cl}(C)$. (Any neighbourhood of x that is included in N must meet exactly the same members of $\mathcal{C}$ as N.) Conversely, $\{C \in \mathcal{C} : x \in \mathrm{cl}(C)\}$ is $\mathcal{C}(N)$.

If $\mathcal{C}$ is a locally finite covering of Ω, then, for each $x \in \Omega$, there is $n(x) \in \mathbb{N} \cup \{0\}$ which is the number of $C \in \mathcal{C}$ such that $x \in \mathrm{cl}(C)$. ****

Definition 11.4. The topological space Ω is *paracompact* if it is Hausdorff and every open covering of Ω has a locally finite open refinement.

I remarked about the definition of compactness that it made no difference whether we required a finite subcovering or a finite refinement. For paracompactness, it is essential to speak of a locally finite *refinement*. For instance, any space with the discrete topology is trivially paracompact in the sense above, since the open covering by singletons is locally finite (indeed discrete) and refines any open covering whatsoever; but, for instance, $\mathbb{N}$ with the discrete topology has an open covering $\{\{1, 2, \ldots, n\} : n \in \mathbb{N}\}$ that cannot have a locally finite subcovering, since any subcovering must have infinitely many members, all containing 1.

As with regularity and normality, Kelley chose to present a slightly different definition that has not caught on. In his book a paracompact space is defined to be regular in his sense, but not necessarily Hausdorff. As before, this is by no means silly, because it adds some precision to the statements of the theorems, but it is not particularly useful either.

The form of the definition allows certain modifications.

Lemma 11.5. *Suppose that every open covering of an arbitrary topological space Ω has a* closed *locally finite refinement. Then every open covering of Ω has an open locally finite refinement.*

Proof. Let $\mathcal{U}$ be an open covering of Ω. Let $\mathcal{C}$ be a closed locally finite refinement of $\mathcal{U}$, and, for each $x \in \Omega$, suppose that $N(x)$ is an open neighbourhood of x such that $\{C \in \mathcal{C} : C \cap N(x) \neq \emptyset\}$ is finite. By hypothesis, the covering $\{N(x) : x \in \Omega\}$ also has a closed locally finite refinement $\mathcal{D}$; each member of $\mathcal{D}$ meets only finitely many members of $\mathcal{C}$.

For each $C \in \mathcal{C}$, set

$$V(C) := \Omega \setminus \left(\bigcup_{D\in\mathcal{D}\ \&\ C\cap D=\emptyset} D\right).$$

This is an open set, since the union is closed by 11.3, and it includes C, since the union is disjoint from C. Thus $\{V(C) : C \in \mathcal{C}\}$ is an open covering of Ω. If $F \in \mathcal{D}$ and $F \cap C = \emptyset$, then $F \cap V(C) = \emptyset$; conversely, if $F \cap V(C) = \emptyset$, then $F \cap C = \emptyset$ (as $V(C) \supseteq C$). Thus F can meet only finitely many of the sets $V(C)$ (namely, those for which it meets C).

For each $C \in \mathcal{C}$, choose $U(C) \in \mathcal{U}$ such that $C \subseteq U(C)$. Let

$$W(C) := V(C) \cap U(C).$$

Certainly $C \subseteq W(C)$, so these sets cover Ω. They are open. For any $x \in \Omega$, there is a neighbourhood M of x that meets (and is covered by) only finitely many members of $\mathcal{D}$. Each of these meets only finitely many $V(C)$; so M can meet only finitely many of the sets $W(C)$. Hence $\{W(C) : C \in \mathcal{C}\}$ is the desired open locally finite refinement of $\mathcal{U}$. □

Lemma 11.6. *A compact Hausdorff space is paracompact.* □

Lemma 11.7. *A separable metric space is paracompact.*

Proof. Indeed, this is the result of Dieudonné's argument at 11.1. □

The most famous result of this kind is the theorem of A. H. Stone:

Theorem 11.8. *Any metric space is paracompact.*

The proof is an extremely clever modification of Dieudonné's argument, and proves rather more than you might expect.

Proof. Let $\mathcal{U} = \{U_s : s \in S\}$ be an open covering of Ω, whose metric is d. Well-order S. Now define, inductively over $n \in \mathbb{N}$, a family $\mathcal{V}_n := \{V_{s,n} : s \in S\}$ of subsets of Ω, also indexed by S, as follows. Suppose $\mathcal{V}_m$ has been defined for all indices $m < n$. (If $n = 1$, this is automatically true.) Define $W_m := \bigcup_{t\in S} V_{t,m}$ for each $m < n$.

Given $s \in S$ and $n \in \mathbb{N}$, let $Q(s,n)$ be the set of points y of Ω such that

- *(a)* $y \in U_s$;
- *(b)* $y \notin U_t$ for $t < s$;
- *(c)* for any $m < n$ and for any $t \in S$, $y \notin V_{t,m}$; that is, $y \notin \bigcup_{m<n} W_m$;
- *(d)* $B(y; 3/2^n) \subseteq U_s$.

(Of course if s is the least element of S, condition *(b)* is vacuously satisfied, whilst if $n = 1$, the condition *(c)* is vacuously satisfied.) Then set

$$V_{s,n} := \bigcup_{y\in Q(s,n)} B(y; 2^{-n}). \tag{27}$$

From (27), $V_{s,n}$ is open, and *(d)* ensures that $V_{s,n} \subseteq U_s$. (28)

The family $\mathcal{V}_n$ is discrete. (In fact, it is even "uniformly discrete" with respect to d, although it is not worth defining what this means.) Suppose that $x_1 \in V_{s_1,n}$ and $x_2 \in V_{s_2,n}$, where we may choose the notation so that $s_1 < s_2$. By (27), there are points $y_i \in Q(s_i, n)$ such that $x_i \in B(y_i; 2^{-n})$ for $i = 1, 2$. But then, by *(b)*, $y_2 \notin U_{s_1}$, and, by *(d)*, $B(y_1; 3/2^n) \subseteq U_{s_1}$. Consequently $d(y_1, y_2) \geq 3/2^n$, and

$$d(x_1, x_2) \geq d(y_1, y_2) - d(x_1, y_1) - d(x_2, y_2) > 2^{-n}.$$

So any point of $V_{s_1,n}$ is distant at least 2^{-n} from any point of $V_{s_2,n}$. Any ball of radius 2^{-n-1} can meet at most one of the sets $V_{s,n}$ in the class $\mathcal{V}_n$, which is, therefore, discrete. (29)

Take any $x \in \Omega$. As $\mathcal{U}$ is a covering of Ω, there is a member of $\mathcal{U}$ containing x. As S is well-ordered, there is, in fact, a *least* s such that $x \in U_s$; and, as U_s is then open, there is some $n \in \mathbb{N}$ such that $B(x; 3/2^n) \subseteq U_s$. It is possible that $x \in V_{t,m}$ for some $m < n$ and some $t \in S$, i.e. that $x \in \bigcup_{m<n} W_m$. But, if this is not true, then $x \in Q(s,n)$, and so $B(x; 2^{-n}) \subseteq V_{s,n} \subseteq W_n$. Either way, we have shown that, for some $m \leq n$, $x \in W_m$. Hence, $\mathcal{V} := \bigcup_{n\in\mathbb{N}} \mathcal{V}_n$ is an open covering of Ω. It is a refinement of $\mathcal{U}$, by the remark (28). And we have shown it is σ-discrete.

It remains to show that $\mathcal{V}$ is locally finite. Given $x \in \Omega$, there are t and m such that $x \in V_{t,m}$, and so by (27) there is some $y \in Q(t,m)$ with $d(x,y) < 2^{-m}$. But there is some $k \in \mathbb{N}$ such that $B(x; 2^{-k}) \subseteq V_{t,m}$. Now suppose that $n \geq m + k$.

For any $s \in S$ and for each $y \in Q(s,n)$, *(c)* says that $y \notin V_{t,m}$, so that $d(y,x) \geq 2^{-k}$. But then $B(x; 2^{-m-k}) \cap B(y; 2^{-n}) = \emptyset$, since otherwise

$$d(x,y) < 2^{-m-k} + 2^{-n} \leq 2^{-n+1} \leq 2^{-k}.$$

This holds for each $y \in Q(s,n)$, so by (27) $B(x; 2^{-m-k}) \cap V_{s,n} = \emptyset$, whenever $n \geq m + k$ (irrespective of $s \in S$). Thus $B(x; 2^{-m-k})$ can meet only members of $\mathcal{V}_1, \mathcal{V}_2, \ldots \mathcal{V}_{m+k-1}$. However, (29) tells us that it meets at most one member of each of these classes. So, in fact, it meets at most $m + k - 1$ members of $\mathcal{V}$, and $\mathcal{V}$ is locally finite. □

Corollary 11.9. *If Ω is pseudometrizable (i.e. its topology may be defined by a pseudometric), then any open covering of Ω has a refinement that is simultaneously locally finite and σ-discrete.* ☐

The proof used σ-discreteness as a step towards local finiteness. There was no place where d had to be a metric rather than a pseudometric — only the triangle inequality was used. However, we cannot simply say that any pseudometric space is paracompact, since our definition of paracompactness requires Ω to be T_2. Kelley's modified definition would have an advantage here.

Corollary 11.10. *The topology of any metrizable space has a σ-discrete base.*

Proof. Take a metric d defining the topology. For $n \in \mathbb{N}$, let

$$\mathcal{U}^{(n)} := \{B(x; \tfrac{1}{n}) : x \in \Omega\},$$

which is an open covering of Ω. By 11.8, there is a σ-discrete (and locally finite) open refinement $\mathcal{V}^{(n)}$ of this covering. The family $\bigcup_{n=1}^{\infty} \mathcal{V}^{(n)}$ is also σ-discrete, and it is a base for the topology. ***** ☐

It is obvious that a σ-discrete family is σ-locally finite as well.

Lemma 11.11. *Let A and B be disjoint closed subsets of a paracompact space Ω. Suppose that, for each $x \in B$, there are disjoint open sets U_x, V_x such that $A \subseteq U_x$ and $x \in V_x$. Then there are disjoint open sets U, V such that $A \subseteq U$ and $B \subseteq V$.*

Proof. Since $\{\Omega \setminus B\} \cup \{V_x : x \in B\}$ is an open covering of Ω, it has an open locally finite refinement $\mathcal{W}$. Let $\mathcal{W}_0 := \{W \in \mathcal{W} : W \cap B \neq \emptyset\}$. Then clearly

$$B \subseteq \bigcup_{W \in \mathcal{W}_0} W. \tag{30}$$

However, if $W \in \mathcal{W}_0$, it must be included in V_x for some $x \in B$ (because it can't be included in $\Omega \setminus B$). Thus $U_x \cap W = \emptyset$, and $U_x \cap \mathrm{cl}(W) = \emptyset$, and $A \cap \mathrm{cl}(W) = \emptyset$. It follows that $A \cap \left(\bigcup_{W \in \mathcal{W}_0} \mathrm{cl}(W)\right) = \emptyset$. By 11.3,

$$C := \bigcup_{W \in \mathcal{W}_0} \mathrm{cl}(W) = \mathrm{cl}\Big(\bigcup_{W \in \mathcal{W}_0} W\Big)$$

is closed. Thus $A \subseteq \Omega \setminus C$, and take $U := \Omega \setminus C$. On the other hand, by (30), $B \subseteq V := \bigcup_{W \in \mathcal{W}_0} W$. As $C \supseteq V$, $U \cap V = \emptyset$. ☐

Theorem 11.12. *Any paracompact space Ω is normal.*

Proof. Given disjoint nonempty closed sets K and L in Ω, and $k \in K$, take in 11.11 $A := \{k\}$ and $B := L$. The hypothesis of 11.11 is satisfied because Ω is Hausdorff. Thus there are disjoint open sets U_k, V_k such that $k \in U_k$, $L \subseteq V_k$. (Thus Ω is regular.) But now we may apply 11.11 again, this time with $A := L$ and $B := K$ (and the symbols "U","V" interchanged!), and the result follows. ☐

Lemma 11.13. *Let Ω be any topological space, and let $\mathcal{U}$ be a σ-locally finite open covering of Ω. Then $\mathcal{U}$ has a locally finite refinement, not necessarily open.*

Proof. Suppose that $\mathcal{U} = \bigcup_{n=1}^{\infty} \mathcal{U}_n$, where each $\mathcal{U}_n$ is a locally finite class. Define $\mathcal{U}'_n := \mathcal{U}_n \setminus (\bigcup_{k<n} \mathcal{U}_k)$ for each n; then the classes $\mathcal{U}'_n$ are disjoint (i.e. a member of $\mathcal{U}$ belongs to only one of the subclasses $\mathcal{U}'_n$), but, as $\mathcal{U}'_n \subseteq \mathcal{U}_n$ for each n, $\mathcal{U}'_n$ is still locally finite; furthermore, $\mathcal{U} = \bigcup_{n=1}^{\infty} \mathcal{U}'_n$ still. For each $U \in \mathcal{U}'_n$, define

$$U^{\sim} := U \setminus \left(\bigcup_{k<n} \bigcup_{V \in \mathcal{U}'_k} V \right). \tag{31}$$

Then $\{U^{\sim} : U \in \mathcal{U}\}$ is a covering of Ω. Indeed, given $x \in \Omega$, there is a least n such that $x \in \bigcup_{V \in \mathcal{U}'_n} V$, and then, for any $U \in \mathcal{U}'_n$ such that $x \in U$, $x \in U^{\sim}$ too.

However, for each $k \in \mathbb{N}$, there is some neighbourhood $M_k \in \mathfrak{N}(x)$ that meets only finitely many members of $\mathcal{U}'_k$. Thus

$$M := U \cap M_1 \cap M_2 \cap \cdots \cap M_n$$

is a neighbourhood of x in Ω that meets only finitely many members V of $\bigcup_{k \le n} \mathcal{U}'_k = \bigcup_{k \le n} \mathcal{U}_k$, and consequently only finitely many of the corresponding smaller sets $V^{\sim}$. (31) shows that $M \subseteq U \in \mathcal{U}_n$ meets no sets $W^{\sim}$ if $W \in \mathcal{U}'_m$ for $m > n$. This proves that $\{U^{\sim} : U \in \mathcal{U}\}$ is locally finite in Ω. (But $U^{\sim}$ is not usually open.) $\square$

Lemma 11.14. *Suppose that Ω is regular, and that every open covering of Ω has a locally finite refinement (consisting of arbitrary sets). Then any open covering has a closed locally finite refinement.*

Proof. Let $\mathcal{U}$ be an open covering of Ω. For each $x \in \Omega$ there is some $U \in \mathcal{U}$ with $x \in U$. By regularity of Ω, there is some open set $V(x)$ such that $x \in V(x)$ and $\mathrm{cl}(V(x)) \subseteq U$. Thus, $\{V(x) : x \in \Omega\}$ is an open covering of Ω; by hypothesis, it has a locally finite refinement $\mathcal{S}$ (by arbitrary sets). For each $S \in \mathcal{S}$, choose $U(S) \in \mathcal{U}$ and $x(S) \in U(S)$ such that $S \subseteq V(x(S))$ and $\mathrm{cl}(V(x(S))) \subseteq U(S)$. Now, for each $U \in \mathcal{U}$, let

$$T(U) := \bigcup_{U(S)=U} S,$$
$$\mathrm{cl}(T(U)) = \bigcup_{U(S)=U} \mathrm{cl}(S) \subseteq U.$$

The equality holds by 11.3, because any subclass of $\mathcal{S}$ is locally finite.

Consider the sets $\bigcup_{U(S)=U} S$, for $U \in \mathcal{U}$. The classes $\{S : U(S) = U\}$ partition $\mathcal{S}$ (i.e. for different U, the corresponding classes are disjoint, but every member of $\mathcal{S}$ appears in some class). Consequently, in the first place

$$\bigcup_{U \in \mathcal{U}} T(U) = \bigcup_{S \in \mathcal{S}} S = \Omega,$$

but also, if an open set V meets only finitely many members of $\mathcal{S}$, say $S_1, \ldots, S_k$, and, as a result, amongst their closures meets only $\mathrm{cl}(S_1), \ldots, \mathrm{cl}(S_k)$, then it meets only $T(U(S_1)), \ldots, T(U(S_k))$ amongst the sets $T(U)$, and amongst their closures only $\mathrm{cl}(T(U(S_1))), \ldots, \mathrm{cl}(T(U(S_k)))$. So the sets $\mathrm{cl}(T(U))$ form a locally finite closed covering of Ω. Finally, for each $S \in \mathcal{S}$

$$\mathrm{cl}(S) \subseteq \mathrm{cl}(V(x(S))) \subseteq U(S)\,,$$

whence $\mathrm{cl}(T(U)) \subseteq U$, and $\{\mathrm{cl}(T(U)) : U \in \mathcal{U}\}$ is a closed refinement of $\mathcal{U}$. □

Corollary 11.15. *The closed locally finite refinement $\{cl(T(U)) : U \in \mathcal{U}\}$, indexed by $\mathcal{U}$, has the property that, for each $U \in \mathcal{U}$, $cl(T(U)) \subseteq U$.* □

In other words, the sets of the original open covering $\mathcal{U}$ may be "shrunk" individually to create a closed locally finite refinement.

Remark 11.16. Suppose that Ω is paracompact (and therefore regular), and that $\mathcal{U}$ is an open covering thereof (which may but need not be locally finite). Then, for each $U \in \mathcal{U}$, one may construct an open set $V(U) \subseteq U$ in such a way that $\mathrm{cl}(V(U)) \subseteq U$ for each U and the class $\{V(U) : U \in \mathcal{U}\}$ is a locally finite covering of Ω. Indeed, one may suppose that $\mathcal{S}$, in the above proof, consists of *open* sets, and then the sets $T(U)$ are also open. Thus, in a paracompact space, each set U of an open covering may be shrunk to a smaller open set whose closure is included in U, in such a way that the shrunken sets constitute a locally finite covering. It should be emphasized that this is not directly stated in the original definition of paracompactness, in which the refinement might in principle have many more members than the original covering.

Theorem 11.17. *Let Ω be regular. Then the following properties are equivalent.*

(a) Ω is paracompact.
(b) Every open covering of Ω has a σ-locally finite open refinement.
(c) Every open covering of Ω has a locally finite refinement (by arbitrary sets).
(d) Every open covering of Ω has a closed locally finite refinement.

Proof. (b) $\Longrightarrow$ (c) by 11.13, — it is obvious that a refinement of a refinement is a refinement. (c) $\Longrightarrow$ (d) by 11.14. Then (d) $\Longrightarrow$ (a) by 11.5. Of course (a) $\Longrightarrow$ (b); a locally finite covering is certainly σ-locally finite. □

The sufficiency of property (b) has a valuable consequence. So far we have only had two classes of paracompact spaces: compact spaces and metrizable spaces.

Definition 11.18. A topological space Ω has the Lindelöf property or is Lindelöf if every open covering of Ω has a countable subcovering.

It is clear that it is immaterial whether we speak here of a "countable refinement" (by arbitrary sets) or of a countable subcovering. The definition I have just given is the usual one, but Engelking requires a Lindelöf space to be regular as well. Obviously a compact space is Lindelöf.

Lemma 11.19. *Any second countable space is Lindelöf.* □

Definition 11.20. Ω is σ-compact if it can be expressed as a countable union of compact subsets.

Lemma 11.21. *A σ-compact topological space is Lindelöf.* $\square$

There are interesting spaces that are σ-compact but not second countable. In any case, we get from 11.24 a considerable widening of the class of paracompact spaces.

Proposition 11.22. *A regular Lindelöf space is paracompact.*

Proof. Indeed, a countable subcovering is σ-locally finite. $\square$

Definition 11.23. Let $\mathcal{U}$ be a covering of Ω. If $A \subseteq \Omega$, the star of A with respect to $\mathcal{U}$, $\mathrm{St}(A,\mathcal{U})$, is the union of the members of $\mathcal{U}$ that meet A. If A is a singleton $\{x\}$, we write $\mathrm{St}(x,\mathcal{U})$. Then a covering $\mathcal{V}$ of Ω is a star-refinement of $\mathcal{U}$ if, for each $V \in \mathcal{V}$, there is some $U \in \mathcal{U}$ such that $\mathrm{St}(V,\mathcal{V}) \subseteq U$. It is a pointwise star-refinement if, for every $x \in \Omega$, there is some $U \in \mathcal{U}$ such that $\mathrm{St}(x,\mathcal{V}) \subseteq U$.

Lemma 11.24. *Suppose that an open covering $\mathcal{U}$ of the topological space Ω has a closed locally finite refinement. Then it also has an open pointwise star-refinement.*

Proof. Let $\mathcal{C}$ be a closed locally finite refinement of $\mathcal{U}$. For each $C \in \mathcal{C}$, choose $U(C) \in \mathcal{U}$ such that $C \subseteq U(C)$. For each $x \in \Omega$, let $\mathcal{C}(x) = \{C \in \mathcal{C} : x \in C\}$; this is certainly a finite subset of $\mathcal{C}$, and so $\bigcap_{C\in\mathcal{C}(x)} U(C)$ is open. Of course $\mathcal{C} \setminus \mathcal{C}(x)$ is locally finite, as $\mathcal{C}$ is, so that $\bigcup_{\mathcal{C}\setminus\mathcal{C}(x)} C$ is closed, and

$$V(x) := \left(\bigcap_{C\in\mathcal{C}(x)} U(C)\right) \cap \left(\Omega \setminus \left(\bigcup_{C\in\mathcal{C}\setminus\mathcal{C}(x)} C\right)\right) \tag{32}$$

is open in Ω. But $x \in V(x)$, so that $\{V(x) : x \in \Omega\}$ is an open covering of Ω.

Let $y \in \Omega$, and choose $C_0 \in \mathcal{C}$ such that $y \in C_0$. If $y \in V(x)$, the second factor in (32) implies that necessarily $y \notin \bigcup_{C\in\mathcal{C}\setminus\mathcal{C}(x)} C$, so that $C_0 \in \mathcal{C}(x)$. But then the first factor shows that $V(x) \subseteq U(C_0)$. This shows that $\{V(x)\}$ is a pointwise star-refinement of $\mathcal{U}$. $\square$

The next lemma is purely set-theoretic.

Lemma 11.25. *Let $\mathcal{U}$ be any covering of Ω. If $\mathcal{V}$ is a pointwise star-refinement of $\mathcal{U}$ and $\mathcal{W}$ is a pointwise star-refinement of $\mathcal{V}$. Then $\mathcal{W}$ is a star-refinement of $\mathcal{U}$.*

Proof. Suppose $W \in \mathcal{W}$. For each $x \in W$, choose some $V(x) \in \mathcal{V}$ such that

$$W \subseteq \mathrm{St}(x,\mathcal{W}) \subseteq V(x)\,; \quad \text{then}$$
$$\mathrm{St}(W,\mathcal{W}) = \bigcup_{x\in W} \mathrm{St}(x,\mathcal{W}) \subseteq \bigcup_{x\in W} V(x)\,.$$

On the other hand, take a specific $w \in W$. Then $w \in W \subseteq V(x)$ for each $x \in W$, and *so* $\bigcup_{x\in W} V(x) \subseteq \mathrm{St}(w,\mathcal{V})$. The result follows. $\square$

The next lemma is more substantial, and indeed difficult.

Lemma 11.26. *Suppose that every open covering of the topological space Ω has an open star-refinement. Then every open covering has an open σ-discrete refinement.*

Proof. Let $\mathcal{U}$ be an open covering of Ω. Define a sequence $(\mathcal{U}_n)$ of open coverings, where $\mathcal{U}_0 = \mathcal{U}$ and $\mathcal{U}_{n+1}$ is a star-refinement of $\mathcal{U}_n$ for each $n \geq 0$. For each $U \in \mathcal{U}$ and each $n \geq 1$, define a corresponding subset of U

$$U_n := \{x \in \Omega : (\exists M \in \mathfrak{N}(x)) \ \ \mathrm{St}(M, \mathcal{U}_n) \subseteq U\}\,. \tag{33}$$

Then $\{U_n : U \in \mathcal{U}\}$ is an open covering of Ω for each $n \geq 1$, and refines $\mathcal{U}$.

Well-order $\mathcal{U}$ by $\leq$, and define for each $U \in \mathcal{U}$ and $n \geq 0$ an open set.

$$U^{(n)} := U_n \setminus \left(\mathrm{cl}\left(\bigcup\nolimits_{V<U} V_{n+1}\right)\right)\,. \tag{34}$$

If $W \in \mathcal{U}_{n+1}$, there is some $X \in \mathcal{U}_n$ such that $\mathrm{St}(W, \mathcal{U}_{n+1}) \subseteq X$, by construction. If $x \in W \cap U_n$, where $U \in \mathcal{U}$, then $x \in X$ and, by (33), $X \subseteq \mathrm{St}(x, \mathcal{U}_n) \subseteq U$; thus, $\mathrm{St}(W, \mathcal{U}_{n+1}) \subseteq U$, which entails $W \subseteq U_{n+1}$, again by (33). This has the consequence that

$$\begin{gathered}\text{if } U \in \mathcal{U},\ x \in U_n, \text{ and } y \notin U_{n+1}, \\ \text{there can be no } W \in \mathcal{U}_{n+1} \text{ with } x, y \in W\,.\end{gathered} \tag{35}$$

Now take any $W \in \mathcal{U}_{n+1}$ (for given n). I claim that W meets at most one of the sets $U^{(n)}$ (see (34)) for $U \in \mathcal{U}$. If it meets $U^{(n)}$ and $V^{(n)}$, where $U < V$, there is some $y \in W \cap V^{(n)} \subseteq W \setminus U_{n+1}$, and some $x \in W \cap U^{(n)} \subseteq U_n$; but (35) shows that this is impossible.

Therefore, the family $\{U^{(n)} : U \in \mathcal{U}\}$ is discrete, for $n = 1, 2, \ldots$.

To complete the proof we must show that $\{U^{(n)} : U \in \mathcal{U}\ \&\ n \in \mathbb{N}\}$ is a covering of Ω. Given a point $y \in \Omega$, there is (by well-ordering) a least member U of $\mathcal{U}$ such that $y \in U_n$ for some $n \in \mathbb{N}$. (Recall that, for each n, the sets U_n for $U \in \mathcal{U}$ cover Ω.) For any $V < U$, then $y \notin V_{n+2}$. From (35),

$$\mathrm{St}(y, \mathcal{U}_{n+2}) \cap V_{n+1} = \emptyset\,,$$

and, taking the union, $\mathrm{St}(y, \mathcal{U}_{n+2}) \cap (\bigcup_{V<U} V_{n+1}) = \emptyset$, so that indeed

$$\mathrm{St}(y, \mathcal{U}_{n+2}) \cap \mathrm{cl}\left(\bigcup\nolimits_{V<U} V_{n+1}\right) = \emptyset\,.$$

But this shows that $y \notin \mathrm{cl}(\bigcup_{V<U} V_{n+1})$, and $y \in U^{(n)}$ by (34). $\square$

Theorem 11.27. *Suppose that the topological space Ω is T_1. Then the following conditions are equivalent.*

(a) *Ω is paracompact.*

(b) *Every open covering of Ω has an open pointwise star-refinement.*

(c) *Every open covering of Ω has an open star-refinement.*

(d) *Ω is regular, and every open covering of Ω has an open σ-discrete refinement.*

Proof. (a) $\Longrightarrow$ (b) by 11.17(d) and 11.24. (b) $\Longrightarrow$ (c) by 11.25. (d) $\Longrightarrow$ (a) by 11.17(b). As for (c) $\Longrightarrow$ (d), 11.26 settles everything but regularity.

Let $x \in \Omega$, and let F be a closed set of Ω not containing x. As Ω is T_1, there is an open covering $\{\Omega \setminus F, \Omega \setminus \{x\}\}$ of Ω. Take an open star-refinement $\mathcal{U}$ of this covering; there is some $U \in \mathcal{U}$ such that $x \in U$. But $\mathrm{St}(U, \mathcal{U})$ is not included in $\Omega \setminus \{x\}$; thus $\mathrm{St}(U, \mathcal{U}) \subseteq \Omega \setminus F$. However,

$$\mathrm{St}(U, \mathcal{U}) = \mathrm{St}(\mathrm{cl}(U), \mathcal{U})$$

(any open set meets U only if it meets $\mathrm{cl}(U)$), so that $\mathrm{cl}(U) \subseteq \mathrm{St}(\mathrm{cl}(U), \mathcal{U}) \subseteq \Omega \setminus F$. Hence,

$$x \in U\,, \quad \Omega \setminus \mathrm{cl}(U) \supseteq F\,,$$

which proves the regularity. $\square$

It is possible to say a great deal more about paracompactness, which is a sort of crossroads where many different properties meet, but I shall not go further here.

12. Partitions of unity.

As before, Ω denotes a topological space.

Definition 12.1. Let $ш : \Omega \longrightarrow \mathbb{R}$. The support of $ш$ is

$$\mathrm{supp}(ш) := \mathrm{cl}_\Omega\{x \in \Omega : ш(x) \neq 0\}\,. \tag{36}$$

There are analogous definitions in some other circumstances (the codomain need not be $\mathbb{R}$, for instance), and sometimes the word "support" is used without the assumption that the closure has been taken.

Definition 12.2. A family Ш of functions $\Omega \longrightarrow \mathbb{R}$ is locally finite if the class

$$\mathcal{C} := \{\mathrm{supp}(ш) : ш \in Ш\}$$

is locally finite in the topological space Ω. A locally finite partition of unity in Ω is a locally finite class Ш of functions $\Omega \longrightarrow [0, \infty)$ such that

$$(\forall y \in \Omega) \quad \sum\nolimits_{ш \in Ш} ш(y) = 1\,. \tag{37}$$

Let $\mathcal{U}$ be a covering of Ω. The partition of unity Ш is said to be subordinate to $\mathcal{U}$ if $\mathcal{C}$ is a refinement of $\mathcal{U}$.

In the literature, slightly different definitions may be found. In particular, I think the phrase "partition of unity" as commonly used assumes the requirement that it should be locally finite. Because of 11.3, it is unimportant for this definition whether we take closures in (36) or not.

We shall be concerned here only with partitions of unity consisting of continuous functions, but in many practical applications other classes of functions (more particularly C^∞ functions) have to be used, and the construction cannot then rely simply on Urysohn's lemma as below.

Granted the definition above, there is, for each $x \in \Omega$, a neighbourhood N of x such that $\mathcal{C}(N)$ is finite (see (26)); then, for any $y \in N$, the only nonzero terms of the sum (37) are those for which $\text{supp}(\text{Ш}) \in \mathcal{C}(N)$, so that, in effect, the sum reduces on N to a finite sum. This is the reason for the importance of partitions of unity (and of the idea of local finiteness).

Lemma 12.3. *Let Ω be paracompact. Then, for any open covering $\mathcal{U}$ of Ω, there is a locally finite partition of unity consisting of continuous functions and subordinate to $\mathcal{U}$.*

Proof. $\mathcal{U}$ has a locally finite open refinement $\mathcal{V}$, which (see 11.16) has an open refinement $\{U(V) : V \in \mathcal{V}\}$ such that, for each $V \in \mathcal{V}$, $\text{cl}(U(V)) \subseteq V$. (This refinement is necessarily locally finite too.) In turn, there is a closed refinement $\{C(V) : V \in \mathcal{V}\}$ of $\{U(V) : V \in \mathcal{V}\}$ such that $C(V) \subseteq U(V)$ for each V. By 11.12, Ω is normal; by Urysohn's lemma 7.16, there is for each $V \in \mathcal{V}$ a continuous function

$$f_V : \Omega \longrightarrow [0,1]$$

such that $f_V(C(V)) = \{1\}$, $f_V(\Omega \setminus U(V)) = \{0\}$. It follows that

$$\text{supp}(f_V) \subseteq U(V),$$

and that the sum $\sum_{V \in \mathcal{V}} f_V$ is "locally finite" or, more precisely, "locally finitely nonzero"; this means that, for any $x \in \Omega$, there is a neighbourhood M of x such that at most finitely many of the functions f_V can take a nonzero value at any point of M. Indeed, this will be true if M meets only finitely many of the sets $U(V)$, say $U(V_1), U(V_2), \dots, U(V_m)$. If $y \in M$, then $y \notin U(V)$ for any $V \in \mathcal{V} \setminus \{V_1, V_2, \dots, V_m\}$, and so $f_V(y) = 0$ for such V. Hence, for $x \in M$, the (formal) sum

$$f(y) := \sum\nolimits_{V \in \mathcal{V}} f_V(y)$$

has nonzero terms only corresponding to the indices $V_1, V_2, \dots, V_m$, and so makes sense (being in effect a finite sum) and defines a continuous function on M. This is true on a suitable neighbourhood of any point of Ω; it follows that the function f is defined and continuous on the whole of Ω.

Furthermore, f takes strictly positive values, indeed values not less than 1. For any $x \in \Omega$, there is some $V \in \mathcal{V}$ such that $x \in C(V)$, since $\{C(V) : V \in \mathcal{V}\}$ is a covering of Ω. But then $f(x) \geq f_V(x) = 1$.

It is therefore legitimate to define for each $x \in \Omega$

$$g_V(x) := \frac{f_V(x)}{f(x)}.$$

Then $\text{supp}(g_V) = \text{supp}(f_V) \subseteq U(V)$, so that $\sum_{V \in \mathcal{V}} g_V$ is also locally finite, and, for any $x \in \Omega$, the sum

$$\sum\nolimits_{V\in\mathcal{V}} g_V(x) = \frac{\sum_{V\in\mathcal{V}} f_V(x)}{f(x)} = 1\,,$$

so that $\{g_V : V \in \mathcal{V}\}$ is indeed a locally finite partition of unity subordinate to $\mathcal{U}$. □

It is possible to use partitions of unity to derive some of the properties of paracompact spaces that we obtained by set-theoretic arguments.

13 Uniform spaces

The theory of topological spaces began with metric spaces, and has been very successful in generalizing ideas connected with continuity. But there are other concepts that arise in metric spaces and are lost in the passage to topologies. Perhaps the most obvious candidate for generalization is uniform continuity. If $f : \Omega \longrightarrow \Psi$ is a mapping between metric spaces, it is *uniformly continuous* if

$$(\forall\epsilon > 0)(\exists\delta > 0) \quad d_\Omega(x, y) < \delta \Longrightarrow d_\Psi(f(x), f(y)) < \epsilon\,;$$

that is, the "δ" that, for a given ϵ, is demanded by the definition of continuity at a specific point x, can be chosen so as to work for all points of Ω simultaneously. This is in principle a much stronger requirement than simple continuity: for instance, if $\mathbb{R}$ carries the standard metric, the function $x \mapsto x^3 + x : \mathbb{R} \longrightarrow \mathbb{R}$ is continuous but not uniformly continuous, since

$$|f(x) - f(y)| = |x - y|(1 + x^2 + xy + y^2)\,.$$

However small we make $|x - y|$, this will be arbitrarily large if $|x^2 + xy + y^2|$ is large enough. But if we give the domain the metric

$$\rho(x, y) := |(x^3 + x) - (y^3 + y)|$$

(which also defines the Euclidean topology), whilst keeping the Euclidean metric on the codomain, then f is uniformly continuous. This shows that uniform continuity is dependent on the metrics used, even when they define the same topology.

The difficulty, put crudely, is that continuity depends on measuring "closeness to a given point", whilst uniform continuity requires a measure of "closeness" defined for pairs of points located anywhere in the space. This suggests that, instead of a metric, one need only have a notion of "closeness". We shall see that "uniform" concepts may be defined in a non-metrizable topological group, for instance.

The definition of a *uniformity* was invented by Weil in 1938, but, as with topologies, it has been reformulated in several ways. I shan't attempt to list them. It is in fact curious that the concept has not had more attention — but it may just be that I have not noticed, or that, in truth, there is not too much to say. Even so, I shall not try to discuss uniformities in any depth.

We require a few preliminary remarks. If R, S are subsets of $\Omega \times \Omega$, they are relations on Ω and so we can construct their composite

$$R \circ S := \{(x,z) \in \Omega \times \Omega : \\ (\exists y \in \Omega)\ (x,y) \in R\ \&\ (y,z) \in S\}$$

and the inverse relations such as

$$R^{-1} := \{(y,x) : (x,y) \in R\}\,.$$

R is *symmetric* if $R = R^{-1}$. In general, $(R^{-1})^{-1} = R$.

It will be convenient to define, for any $A \in \mathcal{P}(\Omega)$,

$$R[A] := \{y \in \Omega : (\exists x \in A)\ (x,y) \in R\}\,,$$

and then $(R \circ S)[A] = R[S[A]]$ trivially. When $A = \{x\}$, we can write $R[x]$ as an abbreviation for $R[\{x\}]$.

The identity relation on Ω is the *diagonal* Δ in $\Omega \times \Omega$:

$$\Delta := \{(x,x) : x \in \Omega\}\,.$$

Definition 13.1. Let Ω be a set. A *uniformity* in Ω is a filter $\mathcal{U}$ in the class of subsets of $\Omega \times \Omega$ that satisfies the following additional properties:

(a) for each $U \in \mathcal{U}$, $\Delta \subseteq U$;

(b) if $U \in \mathcal{U}$, then $U^{-1} \in \mathcal{U}$;

(c) if $U \in \mathcal{U}$, there is some $V \in \mathcal{U}$ such that $V \circ V \subseteq U$.

The members U of the uniformity $\mathcal{U}$ are sometimes called *vicinities* or *entourages* of the uniformity. As with topological spaces, the pair $(\Omega, \mathcal{U})$ of a set and a uniformity thereon is called a *uniform space*, and explicit mention of the uniformity may be suppressed when no ambiguity is possible.

Remark 13.2. In 13.1*(c)*, $V \subseteq U$ (because of 13.1*(a)*).

These properties of a uniformity are abstracted from those of the sets $\{(x,y) : d(x,y) < \epsilon\}$ in $\Omega \times \Omega$ when (Ω, d) is a metric space. As ϵ varies over positive values, these sets form a filter form a filter *base* for which *(a)* holds, since $d(x,x) = 0$; *(b)* holds since $d(x,y) = d(y,x)$; and *(c)* holds since, for any $\epsilon > 0$,

$$\{(x,y) : d(x,y) < \tfrac{1}{2}\epsilon\} \circ \{(x,y) : d(x,y) < \tfrac{1}{2}\epsilon\} \\ \subseteq \{(x,y) : d(x,y) < \epsilon\}\,.$$

In general, the symmetric vicinities of a uniformity will constitute a base (i.e. a filter base) for the uniformity, since, if $U \in \mathcal{U}$, $U \cap U^{-1} \in \mathcal{U}$ by *(b)* and is symmetric.

Definition 13.3. Let $\mathcal{U}_1, \mathcal{U}_2$ be uniformities on Ω. We say that $\mathcal{U}_1$ is *finer* or *stronger* than $\mathcal{U}_2$, or that $\mathcal{U}_2$ is *coarser* or *weaker* than $\mathcal{U}_1$, if $\mathcal{U}_2 \subseteq \mathcal{U}_1$.

There is clearly a finest possible uniformity on Ω; a set $U \subseteq \Omega \times \Omega$ is a vicinity of this uniformity if $\Delta \subseteq U$. The conditions of 13.1 are trivially obvious — in *(c)*, one may take $V = \Delta$. This may be called the *discrete uniformity*. There is also a coarsest possible uniformity; since $\mathcal{U}$ is a filter and nonempty, necessarily $\Omega \times \Omega \in \mathcal{U}$, and the coarsest possible uniformity, which might be called the *indiscrete uniformity*, must be $\{\Omega \times \Omega\}$. It is easy to see that it is a uniformity.

Definition 13.4. A subclass $\mathcal{S}$ of the uniformity $\mathcal{U}$ is a *subbase* for $\mathcal{U}$ if the finite intersections of the members of $\mathcal{S}$ form a base for $\mathcal{U}$.

Definition 13.5. Let $\mathcal{U}$ be a uniformity on Ω. The topology $\mathcal{G}(\mathcal{U})$ on Ω induced by $\mathcal{U}$, or the *uniform topology*, is defined by

$$\mathcal{G}(\mathcal{U}) := \{A \in \mathcal{P}(\Omega) : \underbrace{(\forall a \in A)(\exists U \in \mathcal{U}) \;\; U[a] \subseteq A}_{\text{(P)}}\}\,. \tag{38}$$

Lemma 13.6. *The class $\mathcal{G}(\mathcal{U})$ of subsets of Ω is a topology on Ω.*

Proof. Clearly $\emptyset \in \mathcal{G}(\mathcal{U})$ (the condition (P) being vacuously satisfied) and ((P) being automatic) $\Omega \in \mathcal{G}(\mathcal{U})$. Let $\mathcal{V}$ be any subclass of $\mathcal{G}(\mathcal{U})$. Then, if $x \in \bigcup_{V \in \mathcal{V}} V$, there is some $V_0 \in \mathcal{V}$ with $x \in V_0$, and there is some $U_0 \in \mathcal{U}$ such that $U_0[x] \subseteq V_0$. Then $U_0[x] \subseteq \bigcup_{V \in \mathcal{V}} V$. This proves that the union satisfies (P).

If $V_1, V_2 \in \mathcal{G}(\mathcal{U})$, and $x \in V_1 \cap V_2$, then there exist $U_1, U_2 \in \mathcal{U}$ such that

$$U_1[x] \subseteq V_1\,, \quad U_2[x] \subseteq V_2\,.$$

But $U_1 \cap U_2 \in \mathcal{G}(\mathcal{U})$, and $(U_1 \cap U_2)[x] \subseteq U_1[x] \cap U_2[x] \subseteq V_1 \cap V_2$. This shows that $V_1 \cap V_2$ satisfies (P). Hence, $\mathcal{G}(\mathcal{U})$ is a topology. ☐

It is slightly surprising that this lemma requires so little of the class $\mathcal{U}$ (really, just that, if $U_1, U_2 \in \mathcal{U}$, there exists $U_3 \in \mathcal{U}$ such that $U_3 \subseteq U_1 \cap U_2$). But the other properties of $\mathcal{U}$ are needed to ensure that the topology and the uniformity are related in a reasonable way. It is clear that the discrete uniformity defines the discrete topology and the indiscrete uniformity defines the indiscrete topology.

Lemma 13.7. *Let $\mathcal{U}$ be a uniformity on Ω, and let $\mathcal{G}(\mathcal{U})$ be the induced topology.*

(a) *The interior of $A \in \mathcal{P}(\Omega)$ with respect to $\mathcal{G}(\mathcal{U})$ is*

$$B(A) := \{y \in A : (\exists U \in \mathcal{U}) \;\; U[y] \subseteq A\}\,.$$

(b) *For each $x \in \Omega$, $\{U[x] : U \in \mathcal{U}\}$ is a base of neighbourhoods at x in the topology $\mathcal{G}(\mathcal{U})$.*

(c) *Each vicinity $U \in \mathcal{U}$ is a neighbourhood of Δ in $\Omega \times \Omega$ with respect to the product topology induced by $\mathcal{G}(\mathcal{U})$.*

[Recall that "U is a neighbourhood of Δ" means "U includes an open set that includes Δ". I shall write "int" to denote the interior with respect to $\mathcal{G}(\mathcal{U})$.]

Proof. Suppose $A \in \mathcal{P}(\Omega)$. If $y \in \text{int}(A)$, there is an open set O with $y \in O \subseteq A$; hence, by the definition of $\mathcal{G}(\mathcal{U})$, there is some $U \in \mathcal{U}$ such that $U[y] \subseteq O \subseteq A$, and so $y \in B(A)$. Thus $\text{int}(A) \subseteq B(A)$.

Conversely, suppose $y \in B(A)$, $U[y] \subseteq A$ for some $U \in \mathcal{U}$. By 13.1*(c)*, there is some $V \in \mathcal{U}$ such that $V \circ V \subseteq U$. For any $z \in V[y]$, $V[z] \subseteq U[y] \subseteq A$.

I *think* this is the first place where we really have to appeal to 13.1*(a)*. We wish to conclude from the last paragraph that $z \in B(A)$; but, although $V[z] \subseteq A$, we should not know that $z \in A$ without 13.1*(a)*, which tells us $z \in V[z]$. That being so, however (for any $z \in V[y]$), $V[y] \subseteq B(A)$. As y was any element of $B(A)$, $B(A) \in \mathcal{G}(\mathcal{U})$, from (38). $B(A) \subseteq A$ by definition, so $B(A) \subseteq \text{int}(A)$. We have now proved the inclusion in both directions. Hence, $B(A) = \text{int}(A)$, which is *(a)*.

In the same way, if $V \circ V \subseteq U \in \mathcal{U}$, then $x \in V[x] \subseteq B(U[x])$. So $U[x]$ is a neighbourhood of x. If M is any neighbourhood of x, there is an open set containing x and included in M, and by (38) there is some $W \in \mathcal{U}$ such that $W[x] \subseteq M$. So the sets $U[x]$ form a base of neighbourhoods at x. This is *(b)*.

Again, given a vicinity $U \in \mathcal{U}$, take $V \in \mathcal{U}$ such that $V \circ V \subseteq U$. For any $(x, x) \in \Delta$, $V^{-1}[x] \times V^{-1}[x]$ is a neighbourhood of (x, x) in the product topology (by 13.1*(b)*). But, if $(y, z) \in V^{-1}[x] \times V[x]$, $(x, y) \in V^{-1}$ and $(x, z) \in V$, so that $(y, x) \in V$ and $(x, z) \in V$, $(y, z) \in V \circ V \subseteq U$. Each point of Δ has a neighbourhood included in U. This proves *(c)*. □

Definition 13.8. The uniformity $\mathcal{U}$ on Ω is *separated* if $\bigcap_{U \in \mathcal{U}} U = \Delta$.

Lemma 13.9. *The uniformity $\mathcal{U}$ is separated if and only if the topology $\mathcal{G}(\mathcal{U})$ is* T_0. *But also, if $\mathcal{U}$ is separated, then $\mathcal{G}(\mathcal{U})$ is Hausdorff.*

Proof. Let $(x, y) \in (\Omega \times \Omega) \setminus \Delta$. If $\mathcal{G}(\mathcal{U})$ is T_0, either there is some open set containing x but not y, or vice versa. In the former case, by 13.7, there is some $U \in \mathcal{U}$ such that $y \notin U[x]$, that is, $(x, y) \notin U$. In the latter case, there is $V \in \mathcal{U}$ such that $x \notin V[y]$ or $(y, x) \notin V$; but then $(x, y) \notin V^{-1}$. In either case,

$$(x, y) \notin \bigcap_{U \in \mathcal{U}} U .$$

This holds for any $(x, y) \notin \Delta$, and so $\mathcal{U}$ is separated. The argument reverses.

Suppose $\mathcal{U}$ is separated and $x \neq y$ in Ω. There exists $U \in \mathcal{U}$ with $(x, y) \notin U$, and there exists $V \in \mathcal{U}$ such that $V^{-1} = V$ and $V \circ V \subseteq U$. If $z \in V[x] \cap V[y]$, then $(y, z) \in V$, so $(z, y) \in V$ too, and $(x, z) \in V$; hence $(x, y) \in V \circ V \subseteq U$, and this is denied. So in fact $V[x] \cap V[y] = \emptyset$, and $V[x], V[y]$ are disjoint neighbourhoods of x and y. □

I motivated the idea of a uniformity as a generalization of a metric. But there are other important ways for uniformities to arise, and I want briefly to point out one of them.

Definition 13.10. A *topological group* is a pair $(G, \mathcal{G})$ of a group G and a topology $\mathcal{G}$ on G such that (all) the group operations are continuous with respect to $\mathcal{G}$ and the product topology on $G \times G$. In brief, this means that the mapping

$$(x, y) \mapsto xy^{-1} : G \times G \longrightarrow G$$

is continuous when $G \times G$ is given the product topology.

There are many important examples of topological groups; for instance, the general linear group $GL(n; \mathbb{R})$ of all real invertible $n \times n$ matrices, furnished with the subspace topology it inherits from $\mathbb{R}^{n^2}$, or the "classical groups" like $O(n)$ and $U(n)$.

If G is a topological group, it has two natural uniformities: the *left uniformity* $\mathcal{U}_L$ and the *right uniformity* $\mathcal{U}_R$. Write e for the identity element of G. One customarily writes, if $A \subseteq G$ and $x \in G$, xA for $\{xa : a \in A\}$ and so on.

Definition 13.11. $\mathcal{U}_L := \{U \subseteq G \times G : \bigcap_{g \in G} g^{-1} U[g] \in \mathfrak{N}(e)\}$,
$\mathcal{U}_R := \{U \subseteq G \times G : \bigcap_{g \in G} U[g] g^{-1} \in \mathfrak{N}(e)\}$.

Remark 13.12. It may be clearer to define these uniformities by saying, for instance, that a *base* for $\mathcal{U}_L$ is given by the sets $\widehat{V}_L := \{(x,y) : x^{-1}y \in V\}$ as V varies over $\mathfrak{N}(e)$, and a base for $\mathcal{U}_R$ is given by the sets $\widehat{V}_R := \{(x,y) : xy^{-1} \in V\}$ for $V \in \mathfrak{N}(e)$. This formulation shows that the left uniformity has a countable base if and only if G is first countable, and similarly for the right uniformity. ****

Both the left and the right uniformity define the original topology on G. If G is abelian, they coincide. ****

Theorem 13.13. *Let $(\Omega, \mathcal{U})$ be a uniform space, and suppose that $\mathcal{U}$ has a countable base. Then there is a pseudometric d on Ω which defines the uniformity $\mathcal{U}$.*

A result equivalent to the Theorem itself, albeit not in the language of uniformities, was apparently proved by Alexandrov and Urysohn (together) in 1923.

Proof. Suppose $\mathcal{U}$ has a base $(W_n)_{n=1}^{\infty}$. Define $U_0 := \Omega \times \Omega$, $U_1 := W_1 \cap W_1^{-1}$; if $n \in \mathbb{N}$ and $U_n \in \mathcal{U}$ has been defined, choose in the first place $V_{n+1} \in \mathcal{U}$ so that

$$V_{n+1} \circ V_{n+1} \circ V_{n+1} \subseteq U_n ,$$

using 13.1*(c)*, and then set $U_{n+1} := (V_{n+1} \cap W_{n+1}) \cap (V_{n+1} \cap W_{n+1})^{-1}$. In this way (U_n) is a base for $\mathcal{U}$ consisting of symmetric vicinities, and such that, for each $n \geq 0$,

$$U_{n+1} \subseteq U_{n+1} \circ U_{n+1} \subseteq U_{n+1} \circ U_{n+1} \circ U_{n+1} \subseteq U_n . \tag{39}$$

Define $\phi : \Omega \times \Omega \longrightarrow [0,1]$ by setting

$$\phi(x,y) := \begin{cases} 2^{-n} & \text{if } (x,y) \in U_{n-1} \setminus U_n \text{ for some } n \in \mathbb{N}, \\ 0 & \text{if } (x,y) \in \bigcap_{n=0}^{\infty} U_n . \end{cases} \tag{40}$$

Clearly $0 \leq \phi(x,y) = \phi(y,x) \leq \frac{1}{2}$ for all $x, y \in \Omega$. From this function, which is a "first approximation to the distance", we construct a pseudometric by an argument which is not unfamiliar in other contexts. For any finite sequence $(a_0, a_1, \dots, a_k)$ in Ω, we define its "length according to ϕ"

$$L(a_0, a_1, \dots, a_k) := \sum_{i=1}^{k} \phi(a_{i-1}, a_i) .$$

Then $L(a_0, a_1, \dots, a_k) = L(a_k, a_{k-1}, \dots, a_0)$. For any $(x,y) \in \Omega \times \Omega$, let

$$d(x,y) := \inf\{L(a_0, a_1, \dots, a_k) : x = a_0 \ \& \ y = a_k \ \& \ k \in \mathbb{N}\} , \tag{41}$$

it being understood that, if $k > 1$, then $a_1, \dots, a_{k-1} \in \Omega$ are otherwise unrestricted. The infimum is defined (as the set is non-empty and bounded below by 0 in $\mathbb{R}$). It is clear that $d(x,y) = d(y,x)$, and that, for any $x, y, z \in \Omega$,

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

since a sequence used in defining $d(x,y)$ and one used for $d(y,z)$ may be concatenated to form a sequence for $d(x,z)$. Thus, d is a pseudometric; and (from the sequences of length 2) $d(x,y) \leq \phi(x,y)$, so that

$$U_n \subseteq \{(x,y) : d(x,y) < 2^{-n}\} \tag{42}$$

But d might in principle be identically zero.

I wish to prove that

$$\text{if } L(a_0, a_1, \ldots, a_k) < 2^{-n}\text{, where } k, n \in \mathbb{N}\text{, then } \phi(a_0, a_k) < 2^{-n+1}\,. \tag{43}$$

This is true (for all n) when $k = 1$. Suppose it is true for all n and $1 \leq k \leq m \in \mathbb{N}$. Now suppose that $0 < L(a_0, a_1, \ldots, a_{m+1}) < 2^{-n}$, where n is the largest possible integer for which the inequality holds. There is a largest k (with $0 \leq k \leq m+1$) such that $L(a_0, a_1, \ldots, a_k) < 2^{-n-1}$; in fact $k < m+1$, as otherwise we could have taken $n+1$ instead of n. Then, by the inductive hypothesis, $\phi(a_0, a_k) < 2^{-n}$. But also, as k is the largest possible,

$$2^{-n} > L(a_0, \ldots, a_{k+1}) \geq 2^{-n-1}\,,$$

and so $0 \leq L(a_{k+1}, \ldots, a_{m+1}) < 2^{-n-1}$. Again by the inductive hypothesis, as $m - k \leq m$, $\phi(a_{k+1}, a_{m+1}) < 2^{-n}$. Certainly $\phi(a_k, a_{k+1}) < 2^{-n}$ also (as it is a term in the sum $L(a_0, \ldots, a_{m+1})$).

From the original definition of ϕ at (40), $(a_0, a_k), (a_k, a_{k+1}), (a_{k+1}, a_{m+1}) \in U_n$. From (39), $(a_0, a_{m+1}) \in U_{n-1}$, and so $\phi(a_0, a_{m+1}) < 2^{-n+1}$.

This proves (43) for $k = m+1$, unless $L(a_0, a_1, \ldots, a_{m+1}) = 0$. But in that case $\phi(a_i, a_{i+1}) = 0$ for $0 \leq i \leq m$, so that $(a_i, a_{i+1}) \in U_{q+m}$ for any $q \in \mathbb{N}$. Thus, applying (39) inductively, we have for *any* q

$$(a_0, a_2) \in U_{q+m-1}\,, \ldots, \quad (a_0, a_{i+1}) \in U_{q+m-i}\,, \quad (a_0, a_{m+1}) \in U_q\,,$$

and $\phi(a_0, a_{m+1}) = 0$ by (40). This completes the inductive proof of (43).

Now suppose $d(x,y) < 2^{-n-2}$, where $n \in \mathbb{N}$. By definition (41), there is some sequence $(a_0, a_1, \ldots, a_k)$ with $x = a_0$ and $a_k = y$ and $L(a_0, a_1, \ldots, a_k) < 2^{-n-2}$. But then $\phi(x,y) = \phi(a_0, a_k) < 2^{-n-1}$, by (43), and $(x,y) \in U_n$ by (40). With (42), this shows that

$$\{(x,y) : d(x,y) < 2^{-n-2}\} \subseteq U_n \subseteq \{(x,y) : d(x,y) < 2^{-n}\}\,,$$

and the uniformity defined by the pseudometric d is the same as $\mathcal{U}$. □

It is clear that, conversely, a uniformity defined by a pseudometric will have a countable base.

The pseudometric d will be a metric if and only $\bigcap_{n=1}^{\infty} U_n = \Delta$. This is equivalent to saying that the uniformity is separated. Thus the uniformity is *metrizable* if and only if it is separated and has a countable base.

Essentially the same proof (it is necessary, however, to use the specific bases for $\mathcal{U}_L, \mathcal{U}_R$ suggested at 13.12) yields in the case of topological groups a result sometimes called the Birkhoff-Kakutani theorem, which they proved (separately) in 1936.

Theorem 13.14. *Let G be a topological group. It is metrizable if and only if it possesses a countable base of neighbourhoods of the identity. In that case there are a left-invariant metric and a right-invariant metric that both define the topology.* □

Thus, for instance, $GL(n;\mathbb{R})$ has a *left-invariant* metric defining the topology which was originally constructed from a non-invariant metric. ****

Remark 13.15. The construction of d in 13.13 can be performed whenever one has a sequence (U_n) of symmetric vicinities satisfying (39). But *any* symmetric vicinity $U \in \mathcal{U}$ may be the "U_1" term of such a sequence, from 13.1. Thus there is a pseudometric d_U on Ω such that

$$\mathcal{U} \ni \{(x,y) : d(x,y) < 2^{-3}\} \subseteq U \subseteq \{(x,y) : d(x,y) < 2^{-1}\},$$

and, for any $\epsilon > 0$, $\{(x,y) : d(x,y) < \epsilon\} \in \mathcal{U}$ (as $\epsilon > 2^{-n}$ for some n). Briefly, the uniformity $\mathcal{U}$ may be defined by a collection $\mathcal{D}$ of pseudometrics, in the sense that the collection of sets $\{(x,y) : d(x,y) < \epsilon\}$ as d varies over $\mathcal{D}$ and ϵ varies over $(0,\infty)$ forms a subbase for $\mathcal{U}$. This is an alternative way of describing uniformities.

Notice that the pseudometric d is continuous as a function on $\Omega \times \Omega$ with respect to the product topology. ****

It is natural to ask whether the topologies induced from uniformities are in any way special. There is a very simple answer.

Theorem 13.16. *The topology $\mathcal{G}$ on Ω is induced by a uniformity if and only if, for any $a \in \Omega$ and any closed set F not containing a, there is a continuous function $f : \Omega \longrightarrow [0,1]$ such that $f(a) = 0$ and $f(F) = \{1\}$.*

[Thus a T_0 topology on Ω is generated by a uniformity if and only if it is completely regular, i.e. Tikhonov. From 13.9, the uniformity must be separated, and the topology must be T_2; but, in fact, we see now that it must be $T_{3\frac{1}{2}}$.]

Proof. Suppose that $\mathcal{G} = \mathcal{G}(\mathcal{U})$. Given a and F as in the statement, there is some $U \in \mathcal{U}$ such that $U[a] \cap F = \emptyset$. According to 13.15, there is a continuous pseudometric d such that $\{(x,y) : d(x,y) < 2^{-3}\} \subseteq U$. Define

$$f(x) := \min\{2^3 d(a,x), 1\};$$

this is continuous, and $f(a) = 0$, whilst $f(y) = 1$ whenever $y \notin U[a]$. Thus $\mathcal{G}$ has the asserted property.

Conversely, for each continuous function $f : \Omega \longrightarrow [0,1]$, define an associated pseudometric $d_f(x,y) := |f(x) - f(y)|$. These pseudometrics define a uniformity $\mathcal{U}$ as at 13.15, and it is a tedious exercise to prove that $\mathcal{G}(\mathcal{U}) = \mathcal{G}$. *** □

§14. Uniform notions for uniformities

Now that we have uniformities, there are "categorical" questions to ask. We can also generalize various concepts that previously seemed to need a metric, and specialize others that only needed a topology. It is to be understood that the topology on a uniform space is always the uniform topology.

Definition 14.1. Let $(\Omega, \mathcal{U})$ and $(\Psi, \mathcal{V})$ be uniform spaces, and $f : \Omega \longrightarrow \Psi$ a mapping. f is *uniformly continuous* (with respect to the given uniformities) if, for every $V \in \mathcal{V}$, $(f \times f)^{-1}(V) \in \mathcal{U}$. [Here $f \times f$ denotes the mapping $\Omega \times \Omega \longrightarrow \Psi \times \Psi : (x, y) \mapsto (f(x), f(y))$. The definition says, in effect, that, for any $V \in \mathcal{V}$, there is some $U \in \mathcal{U}$ such that, if $(x, y) \in U$, then $(f(x), f(y)) \in V$. This clearly does generalize the idea of uniform continuity of mappings between metric spaces.]

Remark 14.2. A uniformly continuous mapping is continuous with respect to the uniform topologies.

Uniform spaces and uniformly continuous mappings between them form a category. The isomorphisms in this category are the *uniform homeomorphisms*, i.e. uniformly continuous mappings that have uniformly continuous inverses. One may ask questions such as those of §5.

Definition 14.3. Let Ω be a set. If, for each $\beta \in B$, $(\Psi_\beta, \mathcal{U}_\beta)$ is a uniform space, and $f_\beta : \Omega \longrightarrow \Psi_\beta$ is a mapping, then a uniformity $\mathcal{U}$ on Ω may be defined by specifying that the sets $(f_\beta \times f_\beta)^{-1}(U_\beta)$, where $\beta \in B$ and $U_\beta \in \mathcal{U}_\beta$ are arbitrary, should form a subbase for $\mathcal{U}$. $\mathcal{U}$ is the *uniformity induced on* Ω *by the mappings* f_β, and is the coarsest uniformity on Ω making each f_β uniformly continuous. ***

Two special cases are of interest. If $(\Psi, \mathcal{V})$ is a uniform space, and $A \subseteq \Psi$, there is a uniformity on A induced by the inclusion $i_A : A \longrightarrow \Psi$. This is the *subspace uniformity* $\mathcal{V}_A$ on A, which is easily seen to be $\{(A \times A) \cap V : V \in \mathcal{V}\}$. $(A, \mathcal{V}_A)$ is a *uniform subspace* of $(\Psi, \mathcal{V})$.

Similarly, if $\Omega := \prod_{\beta \in B} \Psi_\beta$, the projections $\pi_\beta : \Omega \longrightarrow \Psi_\beta$ induce a uniform structure on Ω, the *product uniformity*.

Results to those of §14 are easily derived.

Definition 14.4. Let Ω be a uniform space with uniformity $\mathcal{U}$. A filter $\mathfrak{F}$ on Ω is *Cauchy* if, for any $U \in \mathcal{U}$, there is some $Q \in \mathfrak{F}$ such that $Q \times Q \subseteq U$. A net $(x_d)_{d \in D}$ in Ω is Cauchy if, for any $U \in \mathcal{U}$, there is some $c \in D$ such that, whenever $d, e \geq c$, $(x_d, x_e) \in U$.

Remark 14.5. Of course these two definitions, for filters and for nets, are related as I explained at 6.18 and 6.19, and I shall not bother to present the arguments for nets henceforth.

Lemma 14.6. *If* $\mathfrak{F}$ *is a filter in the uniform space* Ω *and converges to* $x \in \Omega$ *in the uniform topology, then* $\mathfrak{F}$ *is Cauchy.*

Proof. Suppose that $U \in \mathcal{U}$. Choose $V \in \mathcal{U}$ such that $V = V^{-1}$ and $V \circ V \subseteq U$. Then, by definition 6.17, there is some $Q \in \mathfrak{F}$ such that $Q \subseteq V[x]$.

Take any $(a, b) \in Q \times Q$. Then $a \in V[x]$, so $(x, a) \in V$, and consequently $(a, x) \in V^{-1} = V$; and $b \in V[x]$, so $(x, b) \in V$. Hence, $(a, b) \in V \circ V \subseteq U$. This proves that $Q \times Q \subseteq U$. $\square$

Definition 14.7. A subset A of the uniform space $(\Omega, \mathcal{U})$ is *complete* if every Cauchy filter in A (in the subspace uniformity) converges to some point of A.

Remark 14.8. The limit of a convergent filter in a Hausdorff space is (trivially) unique (I did not point this out previously); and, therefore, the limit, if any, of a Cauchy filter in a separated uniform space is also unique. Evidently uniform homeomorphisms preserve completeness of subsets.

Lemma 14.9. *If a subset of a uniform space is compact, it is complete.* □

[For a Cauchy filter then has a refinement that converges, by 8.15, and that is enough to ensure the original filter converges.]

Theorem 14.10. *Let* $(\Omega, \mathcal{U})$ *be a uniform space.*

(a) *There exist a complete uniform space* $(\widehat{\Omega}, \widehat{\mathcal{U}})$ *and a mapping* $i : \Omega \longrightarrow \widehat{\Omega}$ *such that* i *is a uniform homeomorphism of* Ω *with the uniform subspace* $i(\Omega)$ *of* $(\widehat{\Omega}, \widehat{\mathcal{U}})$ *and* $i(\Omega)$ *is dense in* $\widehat{\Omega}$.

(b) *Furthermore, if* $(\Psi, \mathcal{V})$ *is a complete uniform space and* $f : \Omega \longrightarrow \Psi$ *a uniformly continuous mapping, there is a unique continuous mapping* $\widehat{f} : \widehat{\Omega} \longrightarrow \Psi$ *such that* $\widehat{f} \circ i = f$ *;* $\widehat{f}$ *is also uniformly continuous.*

(c) *Finally, if* $j : \Omega \longrightarrow \Omega_1$ *is another mapping of* Ω *into a complete uniform space* Ω_1 *enjoying the property* (b)*, then there is a uniform homeomorphism* $q : \widehat{\Omega} \longrightarrow \Omega_1$ *such that* $qi = j$.

Proof. [Sketch only.] Take the set of all Cauchy filters in Ω. For each $U \in \mathcal{U}$, say $(\mathfrak{F}_1, \mathfrak{F}_2) \in \widehat{U}$ for the Cauchy filters $\mathfrak{F}_1, \mathfrak{F}_2$ if

$$(\exists F_1 \in \mathfrak{F}_1 \ \ \& \ \ \exists F_2 \in \mathfrak{F}_2) \ \ F_1 \times F_2 \subseteq U ,$$

and introduce an equivalence relation by

$$\mathfrak{F}_1 \sim \mathfrak{F}_2 \quad \text{if and only if} \quad (\forall U \in \mathcal{U}) \ \ (\mathfrak{F}_1, \mathfrak{F}_2) \in \widehat{U} .$$

Let $\widehat{\Omega}$ be the set of equivalence classes; the $\widehat{U}$ determine the uniformity $\widehat{\mathcal{U}}$ on $\widehat{\Omega}$, and the mapping $i : \Omega \longrightarrow \widehat{\Omega}$ carries each point into the equivalence class of the principal ultrafilter it generates, which is certainly Cauchy. □

$(\widehat{\Omega}, \widehat{\mathcal{U}})$ is called the *completion* of Ω. (It is almost traditional to treat this theorem as obvious.)

Definition 14.11. Let A be a subset of the uniform space $(\Omega, \mathcal{U})$. A is *precompact* if, for any $U \in \mathcal{U}$, there is a finite subset F of Ω (depending on U) such that $A \subseteq U[F]$.

This is of course a "uniform" notion, requiring vicinities for its statement.

Lemma 14.12. A *is precompact as a subset of* Ω *if and only if it is precompact as a subset of itself in the subspace uniformity.* □

That is, the same sets would be indicated if the subset F were required to be a subset of A and U were required to be a vicinity of the subspace uniformity.

Lemma 14.13. (a) *The image of a precompact set under a uniformly continuous mapping is precompact.*

(b) *A set compact in the uniform topology is precompact.* □

Lemma 14.14. *The subset A of the uniform space Ω is precompact if and only if every ultrafilter in A is Cauchy.*

Proof. We may suppose $A = \Omega$. Let Ω be precompact and $\mathfrak{F}$ an ultrafilter, and suppose $U \in \mathcal{U}$. Choose $V \in \mathcal{U}$ such that $V = V^{-1}$ and $V \circ V \subseteq U$. Then there are points $x_1, x_2, \ldots, x_n \in \Omega$ such that $\bigcup_{i=1}^{n} V[x_i] = \Omega$. $\mathfrak{F}$ being an ultrafilter, one of the sets $V[x_i]$ must belong to $\mathfrak{F}$ (see 8.18*(c)*). But then $V[x_i] \times V[x_i] \subseteq U$, and it follows that $\mathfrak{F}$ is Cauchy.

Suppose, conversely, that Ω is not precompact. Thus, there exists some $U \in \mathcal{U}$ such that, for any finite set F in Ω, $U[F] \neq \Omega$. As $U[F_1] \cup U[F_2] = U[F_1 \cup F_2]$, the complements $\Omega \setminus U[F]$ form a filter base $\mathfrak{B}$ in Ω; let $\mathfrak{F}$ be an ultrafilter including $\mathfrak{B}$. As $\mathfrak{F}$ contains $\Omega \setminus U[x]$ for each $x \in \Omega$, it does not contain $U[x]$. However, if $\mathfrak{F}$ were Cauchy, there would be some $Q \in \mathfrak{F}$ (necessarily nonempty) such that $Q \times Q \subseteq U$, and so, for any $q \in Q$, $Q \subseteq U[q]$. But then $U[q] \in \mathfrak{F}$, which contradicts the construction of $\mathfrak{F}$; and $\mathfrak{F}$ cannot be Cauchy. □

Lemma 14.15. *The subset A of the uniform space Ω is compact if and only if it is complete and precompact.*

Proof. 8.19, 14.9, 14.14. □

Lemma 14.16. *Let $f : \Omega \longrightarrow \Psi$ be a continuous map between uniform spaces. If Ω is compact, f is uniformly continuous.*

Proof. Let V be a vicinity of the uniformity $\mathcal{V}$ on Ψ. Choose $W \in \mathcal{V}$ such that $W = W^{-1}$ and $W \circ W \subseteq V$. For each $x \in \Omega$, $W[f(x)]$ is a neighbourhood of $f(x)$ in Ψ, so $f^{-1}(W[f(x)])$ is a neighbourhood of x in Ω and there exists a vicinity U_x of the uniformity $\mathcal{U}$ on Ω such that $U_x[x] \subseteq f^{-1}(W[f(x)])$. Choose $T_x \in \mathcal{U}$ such that $T_x \circ T_x \subseteq U_x$. As $x \in \mathrm{int}(T_x[x])$ by 13.7*(b)*, $\{\mathrm{int}(T_x[x]) : x \in \Omega\}$ is an open covering of Ω, with a finite subcovering given by $x(1), x(2), \ldots, x(n) \in \Omega$. Let

$$T := T_{x(1)} \cap T_{x(2)} \cap \cdots \cap T_{x(n)} \in \mathcal{U}.$$

Now suppose that $(x, y) \in T$. There is some $x(i)$ such that $x \in T_{x(i)}[x(i)]$, and consequently $x \in U_{x(i)}[x(i)]$ and also

$$y \in T[x] \subseteq T[T_{x(i)}[x(i)]] \subseteq T_{x(i)}[T_{x(i)}[x(i)]] \subseteq U_{x(i)}[x(i)].$$

By the definition of $U_{x(i)}$, $f(y) \in W[f(x(i))]$ and $f(x) \in W[f(x(i))]$. Hence

$$(f(x), f(y)) \in W \circ W^{-1} \subseteq V. \qquad \square$$

Theorem 14.17. *A compact Hausdorff space Ω admits one and only one uniformity defining the topology.*

Proof. Indeed, there is a uniformity defining the topology by 13.16 (since 8.5 assures us that Ω is normal, and so Tikhonov). If two uniformities $\mathcal{U}_1, \mathcal{U}_2$ on Ω both define the topology, the identity is a homeomorphism $(\Omega, \mathcal{U}_1) \longrightarrow (\Omega, \mathcal{U}_2)$ and 14.16 shows that it is uniformly continuous; this shows that $\mathcal{U}_1 \supseteq \mathcal{U}_2$. The reverse inclusion must also hold, by symmetry. $\square$

****