Ultrafilters and Tychonoff's Theorem

(October, 2015 version)

G. Eric Moorhouse, University of Wyoming

Tychonoff's Theorem, often cited as one of the cornerstone results in general topology, states that an arbitrary product of compact topological spaces is compact. In the 1930's, Tychonoff published a proof of this result in the special case of $[0,1]^A$, also stating that the general case could be proved by similar methods. The first proof in the general case was published by Čech in 1937. All proofs rely on the Axiom of Choice in one of its equivalent forms (such as Zorn's Lemma or the Well-Ordering Principle). In fact there are several proofs available of the equivalence of Tychonoff's Theorem and the Axiom of Choice—from either one we can obtain the other. Here we provide one of the most popular proofs available for Tychonoff's Theorem, using ultrafilters. We see this approach as an opportunity to learn also about the larger role of ultrafilters in topology and in the foundations of mathematics, including non-standard analysis. As an appendix, we also include an outline of an alternative proof of Tychonoff's Theorem using a transfinite induction, using the two-factor case $X \times Y$ (previously done in class) for the inductive step. This approach is found in the exercises of the Munkres textbook.

Henri Cartan (1904–2008) | Alfred Tarski (1901–1983) | Andrey Tikhonov (1906–1993)

The proof of existence of nonprincipal ultrafilters was first published by Tarski in 1930; but the concept of ultrafilters is attributed to H. Cartan.

1. Compactness and the Finite Intersection Property

Before proceeding, let us recall that a collection of sets $\mathfrak{S}$ has the **finite intersection property** if $S_1 \cap \cdots \cap S_n \neq \varnothing$ for all $S_1, \ldots, S_n \in \mathfrak{S}$, $n \geqslant 0$. We point out the following characterization of compact sets by the finite intersection property.

1.1 Lemma. Let X be a topological space. The following two conditions are equivalent:

(i) X is compact.
(ii) If $\mathfrak{S}$ is any collection of closed subsets of X with the finite intersection property, then $\bigcap \mathfrak{S} \neq \infty$.

Proof. Let $\mathfrak{S}$ be any collection of closed sets in X with $\bigcap\mathfrak{S} = \varnothing$, i.e. $\bigcup_{C\in\mathfrak{S}}(X\smallsetminus C) = X\smallsetminus\bigcap_{C\in\mathfrak{S}} C = X$. If X is compact then there exist $C_1,\ldots,C_n \in \mathfrak{S}$ such that $X = (X\smallsetminus C_1)\cup\cdots\cup(X\smallsetminus C_n)$, i.e. $C_1\cap\cdots\cap C_n = \varnothing$ and $\mathfrak{S}$ does not have the finite intersection property. This gives (i)$\Rightarrow$(ii); and the converse follows by reversing the steps. $\square$

A collection of **basic closed sets** is a collection $\mathfrak{C}$ of closed sets such that every closed subset of X is an intersection of sets in $\mathfrak{C}$; that is, $\{X\smallsetminus C : C\in\mathfrak{C}\}$ is a collection of basic open sets. In class we proved (i)$\Leftrightarrow$(ii) of the following; and the third equivalence follows by arguments similar to Lemma 1.1.

1.2 Lemma. Let X be a topological space, and let $\mathfrak{C}$ be a collection of basic closed sets for X. The following three conditions are equivalent:

(i) X is compact.

(ii) Every basic open cover of X has a finite subcover.

(iii) If $\mathfrak{S}\subseteq\mathfrak{C}$ has the finite intersection property, then $\bigcap\mathfrak{S}\neq\varnothing$. $\square$

2. Convergence of ultrafilters

Let X be a topological space, and $\mathfrak{U}$ an ultrafilter on X. We say that $\mathfrak{U}$ **converges** to a point $x\in X$, denoted $\mathfrak{U}\searrow x$, if every open neighborhood U of x satisfies $U\in\mathfrak{U}$. If X is discrete, then every convergent ultrafilter is principal: if $\mathfrak{U}\searrow x$ and $\{x\}$ is open, then $\{x\}\in\mathfrak{U}$ so $\mathfrak{U} = \mathfrak{F}_{\{x\}} = \{U\subseteq X : x\in U\}$. In this case, of course, x is the unique point to which $\mathfrak{U}$ converegesǃ.

2.1 Theorem. Let X be a topological space. Then

(a) X is Hausdorff iff every ultrafilter on X converges to *at most* one point.

(b) X is compact iff every ultrafilter on X converges to *at least* one point.

Proof. (a) Let X be a Hausdorff space, and suppose $\mathfrak{U}\searrow x$ and $\mathfrak{U}\searrow y$ for some points $x\neq y$. Let $U,V\subseteq X$ be disjoint open neighborhoods of x and y respectively. Then $U,V\in\mathfrak{U}$ and so $\varnothing = U\cap V\in\mathfrak{U}$, a contradiction.

Conversely, suppose every ultrafilter on X converges to at most one point; and let $x\neq y$ be points in X. Suppose that every open neighborhood of x has nontrivial intersection with every open neighborhood of y. Then the collection

$$\mathfrak{S} = \big\{\text{open } U\subseteq X \,:\, x\in U \text{ or } y\in U\big\}$$

has the finite intersection property, so $\mathfrak{S}\subseteq\mathfrak{U}$ for some ultrafilter $\mathfrak{U}$. If $U,V\subseteq X$ are open neighborhoods of x and y respectively, then $U,V\in\mathfrak{U}$ so $U\cap V\neq\varnothing$.

(b) Let X be a compact space, and suppose $\mathfrak{U}$ is an ultrafilter on X which does not converge to any point of X. Then for every point $x \in X$, we can find an open neighborhood U_x of x such that $U_x \notin \mathfrak{U}$, i.e. the complement $X \smallsetminus U_x \in \mathfrak{U}$. We obtain a collection of closed sets $\{X \smallsetminus U_x : x \in X\}$ with empty intersection since $x \notin X \smallsetminus U_x$. By Lemma 0.1, there exist $x_1, x_2, \dots, x_n \in X$ such that i.e. $(X \smallsetminus U_{x_1}) \cap \cdots \cap (X \smallsetminus U_{x_n}) = \varnothing$. This is impossible since each $X \smallsetminus U_{x_i} \in \mathfrak{U}$.

Conversely, suppose every ultrafilter on X converges; and suppose that X has an open cover $\mathcal{O}$ without any finite subcover. This says that the collection of closed sets

$$\mathfrak{S} = \{X \smallsetminus U : U \in \mathcal{O}\}$$

has the finite intersection property. Extend this to an ultrafilter $\mathfrak{U} \supseteq \mathfrak{S}$. By hypothesis, there exists a point $x \in X$ such that $\mathfrak{U} \searrow x$. The point x is covered by some set $U \in \mathcal{O}$; and so $U \in \mathfrak{U}$. But also $X \smallsetminus U \in \mathfrak{S} \subseteq \mathfrak{U}$, so $\mathfrak{U}$ contains $(X \smallsetminus U) \cap U = \varnothing$, a contradiction. $\square$

We have the following characterization of the topology by ultrafilters:

2.2 Theorem. Let X be a topological space, and let $U \subseteq X$. Then U is open iff $U \in \mathfrak{U}$ for every ultrafilter $\mathfrak{U}$ converging to some point of U.

Proof. If $U \subseteq X$ is open and $\mathfrak{U}$ is an ultrafilter converging to a point $u \in U$, then we must have $U \in \mathfrak{U}$ by definition of the convergence $\mathfrak{U} \searrow u$.

Conversely, let $U \subseteq X$ and suppose $U \in \mathfrak{U}$ for every ultrafilter $\mathfrak{U}$ converging to some point of U. If U is not open, then there exists a point $u \in U$ such that every open neighborhood of u contains at least one point of $X \smallsetminus U$. This would mean that the collection of sets

$$\mathfrak{S} = \{X \smallsetminus U\} \cup \{\text{all open neighborhoods of } u\}$$

has the finite intersection property. Extend this family of sets to an ultrafilter $\mathfrak{U} \supseteq \mathfrak{S}$. By construction, $\mathfrak{U}$ contains every open neighborhood of u and so $\mathfrak{U} \searrow u$. By hypothesis, $U \in \mathfrak{U}$; but then $\varnothing = (X \smallsetminus U) \cap U \in \mathfrak{U}$, a contradiction. So U must in fact be open as claimed. $\square$

3. Pushing forward ultrafilters

Let $f : X \to Y$ be any map of sets, and suppose $\mathfrak{U}$ is an ultrafilter on X. Define $f_*(\mathfrak{U})$ to be the collection of all subsets $V \subseteq Y$ such that $f^{-1}(V) \in \mathfrak{U}$.

3.1 Proposition. $f_*(\mathfrak{U})$ is an ultrafilter on Y.

Proof. Since $f^{-1}(\varnothing) = \varnothing \notin \mathfrak{U}$ and $f^{-1}(Y) = X \in \mathfrak{U}$, we have $\varnothing \notin f_*(\mathfrak{U})$ and $Y \in f_*(\mathfrak{U})$. If $V_1, V_2 \in f_*(\mathfrak{U})$ then $f^{-1}(V_1 \cap V_2) = f^{-1}(V_1) \cap f^{-1}(V_2) \in \mathfrak{U}$ so $V_1 \cap V_2 \in f_*(\mathfrak{U})$. Also if $V_1 \in f_*(\mathfrak{U})$ and $V_1 \subseteq V \subseteq Y$, then $f^{-1}(V) \supseteq f^{-1}(V_1) \in \mathfrak{U}$ so $f^{-1}(V) \in \mathfrak{U}$ and $V \in f_*(\mathfrak{U})$. So $f_*(\mathfrak{U})$ is a filter on Y.

Finally, if $Y = Y_1 \sqcup Y_2$ (our notation for a disjoint union: $Y = Y_1 \cup Y_2$ with $Y_1 \cap Y_2 = \varnothing$) then $X = f^{-1}(Y_1) \sqcup f^{-1}(Y_2)$ so $f^{-1}(Y_i) \in \mathfrak{U}$ for exactly one choice of $i \in \{1, 2\}$, giving $Y_i \in f_*(\mathfrak{U})$. Thus $f_*(\mathfrak{U})$ is in fact an ultrafilter on Y. □

The ultrafilter $f_*(\mathfrak{U})$ is called the **push-forward** of the ultrafilter $\mathfrak{U}$. Note that if $f : X \to Y$ and $g : Y \to Z$ then $(g \circ f)_*(\mathfrak{U}) = g_*(f_*(\mathfrak{U}))$ is an ultrafilter on Z. The rule $(g \circ f)_* = g_* \circ f_*$ is indicative of 'push-forward' maps; in the case of pull-backs, where *superscript* $*$'s are used, the rule is instead $(g \circ f)^* = f^* \circ g^*$. [Remark: Denote by **Set** the category of sets, where morphisms are maps between sets. The map $X \mapsto$ {ultrafilters on X}, $f \mapsto f_*$ is a covariant functor **Set** $\to$ **Set**.]

3.2 Theorem. Let X and Y be topological spaces, and consider an arbitrary map $f : X \to Y$. Then f is continuous iff $f_*(\mathfrak{U}) \searrow f(x)$ in Y for every ultrafilter $\mathfrak{U} \searrow x$ in X.

Proof. First suppose f is continuous, and let $\mathfrak{U}$ be an ultrafilter on X with $\mathfrak{U} \searrow x$. For every open neighborhood $V \subseteq Y$ of $f(x)$, the set $f^{-1}(V) \subseteq X$ is an open neighborhood of x. We have $f^{-1}(V) \in \mathfrak{U}$ and therefore $V \in f_*(\mathfrak{U})$. By definition, $f_*(\mathfrak{U}) \searrow f(x)$.

Conversely, suppose $f_*(\mathfrak{U}) \searrow f(x)$ in Y whenever $\mathfrak{U}$ is an ultrafilter converging to a point x in X; and let $V \subseteq Y$ be open. To show that the preimage $f^{-1}(V) \subseteq X$ is open, we will use Theorem 2.2. Accordingly, let $\mathfrak{U}$ be any ultrafilter in X converging to some point $x \in f^{-1}(V)$. By hypothesis, $f_*(\mathfrak{U}) \searrow f(x)$ in Y; and since $V \subseteq Y$ is an open neighborhood of $f(x)$, we have $V \in f_*(\mathfrak{U})$, i.e. $f^{-1}(V) \in \mathfrak{U}$. By Theorem 2.2, $f^{-1}(V)$ is open. Thus f is continuous. □

The latter result begs for a comparison between ultrafilters and sequences. Recall that for a continuous function $f : X \to Y$, we have $f(x_n) \to f(x)$ in Y whenever $x_n \to x$ in X. The converse does not hold without some additional hypotheses (e.g. X is a metric space). Sequences of the form $x_1, x_2, x_3, \dots$ are limited in that they are countable by defintion. Ultrafilters, which do not suffer from this constraint, suffice to characterize topologies and continuity in general.

4. Product spaces

Let X_α be a family of topological spacesindexed by $\alpha \in A$. The **product space** $X = \prod_{\alpha \in A} X_\alpha$ has points of the form $x = (x_\alpha)_{\alpha \in A}$ where $x_\alpha \in X_\alpha$ for all $\alpha \in A$. We describe below the natural topology on this set, imposed by the individual topologies on the spaces X_α. But first, let us be clear what is meant by a point $(x_\alpha)_\alpha$.

In the case $|A| = n < \infty$, we might as well take $A = \{0, 1, 2, \ldots, n{-}1\}$. To simplify notation, we use ordinal notation where the finite ordinals are defined recursively by $n = \{0, 1, 2, \ldots, n{-}1\}$, and then

$$X \;=\; X_0 \times X_1 \times \cdots \times X_{n-1} \;=\; \big\{(x_0, x_1, \ldots, x_{n-1}) : x_i \in X_i \text{ for all } i \in n\big\}.$$

In the special case where $X_0 = X_1 = \cdots = X_{n-1}$, we obtain

$$X = \underbrace{X_0 \times X_0 \times \cdots \times X_0}_{n \text{ times}} = X_0^n.$$

Every point $x = (x_0, x_1, \ldots, x_{n-1}) \in X$ can be viewed as a function $f : n \to X_0$; simply identify each function $f : n \to X_0$ with its sequence of values $(f(0), f(1), \ldots, f(n{-}1))$. To accommodate arbitrary finite products $\prod_{i\in n} X_i$, we again view every point as a function $f : n \to \bigcup_i X_i$, but with the restriction that $f(i) \in X_i$ for each i. If we omit the 'f', the function mapping $i \mapsto x_i$ is identified simply by its list of values $(x_0, x_1, x_2, \ldots, x_{n-1})$.

Next we proceed to countable products, using $\omega = \{0, 1, 2, \ldots\}$ as index set. Here the Greek letter ω is the smallest infinite ordinal: it is the set of all finite ordinals, i.e. the set of all non-negative integers. (One could use positive integers instead.) Now we have a countable product of sets given by

$$X \;=\; \prod_{i\in\omega} X_i \;=\; X_0 \times X_1 \times X_2 \times \cdots \;=\; \big\{(x_0, x_1, x_2, \ldots) \;:\; x_i \in X_i \text{ for all } i \in \omega\big\}.$$

Here every point $x \in X$ is an infinite sequence, which we identify with the map $i \mapsto x_i$. As special cases (where all component spaces X_i are equal), we have

$$\mathbb{R}^\omega = \mathbb{R} \times \mathbb{R} \times \mathbb{R} \times \cdots,$$

the set of all (countably) infinite sequences of real numbers, and

$$2^\omega = \{0,1\}^\omega = \{0,1\} \times \{0,1\} \times \{0,1\} \times \cdots,$$

the set of all (countably) infinite binary sequences. (As a set, this last example is the Cantor set.)

Now for an index set A of arbitrary cardinality, the general product $X = \prod_{\alpha\in A} X_\alpha$ may be regarded as the set of all functions defined on A, which map $\alpha \mapsto x_\alpha \in X_\alpha$ for each $\alpha \in A$. Points may be denoted $x = (x_\alpha)_{\alpha\in A}$. In the special case where all X_α's are equal to some fixed space X_0, we have

$$X = X_0^A = \{\text{functions } A \to X_0\}.$$

For example, the set $\mathbb{R}^\mathbb{R}$ is simply the set of all functions $\mathbb{R} \to \mathbb{R}$.

Every product space $X = \prod_{\alpha \in A} X_\alpha$ comes naturally equipped with projections onto the various factors; these are the surjections given by

$$\pi : X \to X_\alpha,\ x \mapsto x_\alpha.$$

Thinking of a point $x \in X$ as a function, the αth coordinate $\pi_\alpha(x) = x_\alpha$ is just the value of that function at the input value $\alpha \in A$.

Now the **product topology** on X is the coarsest topology for which each of the projections $\pi_\alpha : X \to X_\alpha$ is continuous. Equivalently, a basis for the topology on X is the collection of all basic open sets of the form

$$\prod_{\alpha \in A} U_\alpha$$

where each $U_\alpha \subseteq X_\alpha$ is open, *and* $U_\alpha = X_\alpha$ for all but finitely many α.

By contrast, the **box topology** on X is the topology having as basis all sets of the form $\prod_{\alpha \in A} U_\alpha$ where $U_\alpha \subseteq X_\alpha$ for all $\alpha \in A$, and no additional restriction. Of course if $|A| < \infty$, then this coincides with the product topology; but in general, the box topology is a refinement of the product topology—much too fine to be useful for most purposes. Unless otherwise specified, it is the product topology that we will take for a general product of topological spaces $X = \prod_\alpha X_\alpha$.

For example, consider the sequence of points $v_1, v_2, v_3, \ldots \in [0,1]^\omega$ given by

$$\begin{aligned}
v_1 &= (0,1,1,1,1,1,\ldots),\\
v_2 &= (0,0,1,1,1,1,\ldots),\\
v_3 &= (0,0,0,1,1,1,\ldots),\\
v_4 &= (0,0,0,0,1,1,\ldots),
\end{aligned}$$

etc. You might hope that v_n converges to the point $(0,0,0,0,0,0,\ldots)$ in $[0,1]^\omega$, and this is certainly true—if we denote $\mathbf{0} = (0,0,0,\ldots)$ then a basic open neighborhood of $\mathbf{0}$ in the product topology has the form

$$U = U_0 \times U_1 \times U_2 \times \cdots \times U_{m-1} \times [0,1] \times [0,1] \times [0,1] \times \cdots$$

where each U_i is an open neighborhood of 0. Since $v_n \in U$ whenever $n \geqslant m$, we have $v_n \to \mathbf{0}$. In the box topology, the sequence v_n does not converge at all; for example consider the neighborhood

$$U' = \left[0, \tfrac{1}{2}\right] \times \left[0, \tfrac{1}{2}\right] \times \left[0, \tfrac{1}{2}\right] \times \left[0, \tfrac{1}{2}\right] \times \cdots,$$

a basic open neighborhood of $\mathbf{0}$ in the box topology. The sequence v_n never gets inside U', no matter how large n is; so v_n fails to converge to $\mathbf{0}$ in the box topology. You might think this has to do with the coordinates of v_n converging 'pointwise' (and not 'uniformly') to 0; but actually the problem is much worse than that. Consider the sequence

$$\begin{aligned} w_1 &= \big(0, 1, 1, 1, 1, 1, \ldots\big), \\ w_2 &= \big(0, 0, \tfrac{1}{2}, \tfrac{1}{2}, \tfrac{1}{2}, \tfrac{1}{2}, \ldots\big), \\ w_3 &= \big(0, 0, 0, \tfrac{1}{4}, \tfrac{1}{4}, \tfrac{1}{4}, \ldots\big), \\ w_4 &= \big(0, 0, 0, 0, \tfrac{1}{8}, \tfrac{1}{8}, \ldots\big), \end{aligned}$$

etc. which converges to $\mathbf{0}$ (in the product topology). The coordinates do converge 'uniformly' to 0, and reasonably fast, yet the sequence w_n never gets inside

$$\big[0, 1\big] \times \big[0, \tfrac{1}{2}\big] \times \big[0, \tfrac{1}{3}\big] \times \big[0, \tfrac{1}{5}\big] \times \big[0, \tfrac{1}{9}\big] \times \cdots,$$

so $w_n \not\to \mathbf{0}$ in the product topology. The box topology contains far too many open sets—it is too close to the discrete topology to be very useful for us.

Generalizing the examples above, we see that a sequence of points in $X = \prod_\alpha X_\alpha$ converges, iff it converges 'coordinatewise':

4.1 Theorem. Let $x_n = (x_{n,\alpha})_\alpha$ be a sequence of points in $X = \prod_\alpha X_\alpha$. (Note that two subscripts are used: $n = 1, 2, 3, \ldots$ indexes the points of the sequence, and $\alpha \in A$ indexes the coordinates of each point.) Also let $a = (a_\alpha)_\alpha \in X$. Then $x_n \to a$ in X, iff $x_{n,\alpha} \to a_\alpha$ for each α, as $n \to \infty$.

Proof. Suppose first that $x_n \to a$. Since $\pi_\alpha : X \to X_\alpha$ is continuous, this implies that $x_{n,\alpha} = \pi(x_n) \to \pi_\alpha(a) = a_\alpha$ for all α.

To prove the converse, it suffices to consider a basic open neighborhood U of $a \in X$. This is a set of the form $U = \prod_\alpha U_\alpha$ where each $U_\alpha \subseteq X_\alpha$ is an open neighborhood of a_α; and $U_\alpha = X_\alpha$ for all $\alpha \neq \alpha_1, \alpha_2, \ldots, \alpha_m$. Under the hypothesis that $x_{n,\alpha} \to a_\alpha$ for all $\alpha \in A$, there exist constants $N_1, \ldots, N_m$ such that for all $i \in \{1, 2, \ldots, m\}$, we have $x_{n,\alpha_i} \in U_{\alpha_i}$ whenever $n > N_i$. Let $N = \max\{N_1, \ldots, N_m\}$; then for all $n > N$ and all $\alpha \in A$, we have $x_{n,\alpha} \in U_\alpha$ whenever $n > N$. For $\alpha \in \{\alpha_1, \ldots, \alpha_m\}$, this follows by the choice of $n > N \geqslant N_i$; and for $\alpha \notin \{\alpha_1, \ldots, \alpha_m\}$, it follows simply because $x_\alpha \in X_\alpha = U_\alpha$. $\square$

Analogously, the product topology can be characterized using ultrafilters:

4.2 Theorem. Let $\mathfrak{U}$ be an ultrafilter on $X = \prod_{\alpha \in A} X_\alpha$, and let $x = (x_\alpha)_\alpha \in X$. Then $\mathfrak{U} \searrow x$ iff $(\pi_\alpha)_*(\mathfrak{U}) \searrow x_\alpha$ in each X_α.

Proof. If $\mathfrak{U} \searrow x$ then by Theorem 3.2, for each $\alpha \in A$ we have $(\pi_\alpha)_*(\mathfrak{U}) \searrow x_\alpha$.

Conversely, suppose that for each $\alpha \in A$, we have $(\pi_\alpha)_*(\mathfrak{U}) \searrow x_\alpha$. Let $U \subseteq X$ be an open neighborhood of x; we must show that $U \in \mathfrak{U}$. As usual, it suffices to consider subbasic open sets of the form $U = \pi_\alpha^{-1}(U_\alpha)$ where $U_\alpha \subseteq X_\alpha$ is an open neighborhood of x_α. In this case $U_\alpha \in (\pi_\alpha)_*(\mathfrak{U})$ and so $U = \pi_\alpha^{-1}(U_\alpha) \in \mathfrak{U}$ as required. Since $\mathfrak{U}$ is closed under finite intersections and taking supersets, the result carries over for a general open neighborhood U of x, giving $\mathfrak{U} \searrow x$. $\square$

5. Tychonoff's Theorem

5.1 Theorem (Tychonoff). If each of the topological spaces X_α is compact, then so is the product space $X = \prod_\alpha X_\alpha$.

For example, $[0,1]^n$ is compact. Also since closed subsets of a compact space are compact, we see that a subset $K \subseteq \mathbb{R}^n$ is compact iff it is closed and bounded (with respect to the usual metric)... the details of this argument were given in class. Note that the product topology is assumed here—if we substitute the box topology, the result fails, yet another example that the box topology is usually not a good choice. For example, consider the subsets of $[0,1]^\omega$ of the form

$$U_0 \times U_1 \times U_2 \times \cdots$$

where each U_i is either $\left[0, \frac{2}{3}\right)$ or $\left(\frac{1}{3}, 1\right]$. There are $2^{\aleph_0}$ such sets, and they cover $[0,1]^\omega$. They are all open in the box topology; but no finite subcollection of these sets suffice to cover $[0,1]^\omega$.

Proof of Theorem 5.1. We use Theorem 2.1(b). Let $\mathfrak{U}$ be an ultrafilter on X. For each α, the push-forward $(\pi_\alpha)_*(\mathfrak{U})$ is an ultrafilter on X_α. By Theorem 2.1, $(\pi_\alpha)_*(\mathfrak{U}) \searrow x_\alpha$ for some point $x_\alpha \in X_\alpha$. By Theorem 4.2, $\mathfrak{U} \searrow x$ where $x = (x_\alpha)_\alpha \in X$. Since every ultrafilter on X converges, Theorem 2.1 shows that X is compact. $\square$

6. Application: Weak-* Topology

Let V be a real vector space with a norm $\|\cdot\| : V \to \mathbb{R}$ satisfying

- $\|v\| \geqslant 0$, and equality holds iff $v = 0$;
- $\|v + w\| \leqslant \|v\| + \|w\|$ for all $v, w \in V$; and
- $\|cv\| = |c|\|v\|$ for all $c \in \mathbb{R}$, $v \in V$.

We call V a **normed vector space**. A **bounded linear functional** on V is a map $f : V \to \mathbb{R}$ such that

- f is linear: $f(av + bw) = af(v) + bf(w)$ for all $a, b \in \mathbb{R}$, $v, w \in V$;
- there exists a real constant $C \geqslant 0$ such that $|f(v)| \leqslant C\|v\|$ for all $v \in V$.

The set of all bounded linear functionals on V is a vector space, denoted V^*. Consider the closed unit ball in V defined by

$$B = \{v \in V \,:\, \|v\| \leqslant 1\}.$$

Let $f \in V^*$. Since the values of f on B are bounded, we may define

$$\|f\| = \sup\{|f(v)| : v \in B\}.$$

This makes V^* also a normed vector space, hence a metric space with distance function $d(f, g) = \|f - g\|$. The unit ball in V^* is

$$B^* = \{f \in V^* \,:\, \|f\| \leqslant 1\} \;=\; \{f \in V^* \,:\, |f(v)| \leqslant \|v\| \text{ for all } v \in V\}.$$

Note that $|f(v)| \leqslant 1$ for all $f \in B^*$ and $v \in B$. Using linearity, every $f \in B^*$ is uniquely determined by its restriction to B; so after identifying f with this restriction, B^* is the set of all functions $B \to [-1, 1]$ such that f is the restriction of a linear functional $V \to \mathbb{R}$.

There are two reasonable topologies to take on V^*. One is the metric topology given by its norm. This topology is often too strong to be useful; for example the unit ball B^* is not compact in this topology, except in the finite-dimensional case. To see this, observe that B^* is covered by open balls of radius $\frac{1}{2}$ and no finite number of these balls suffice to cover B^* (except in the finite-dimensional case).

To rectify this problem, consider V^* as a subset of the product space $\mathbb{R}^V$. Suppose $f, f_1, f_2, f_3, \ldots \in V^*$. By Theorem 4.1, $f_n \to f$ in this topology iff we have pointwise convergence $f_n(v) \to f(v)$ for every $v \in V$. This topology is called the **weak-* topology** on V^*. It is coarser (weaker) than the norm topology (strong topology) defined above. Note that pointwise convergence does not imply $\|f_n - f\| \to 0$, so the convergence $f_n \to f$ does not hold in the norm topology.

6.1 Theorem. The subset $B^* \subseteq V^*$ is closed and compact in the weak-* topology.

To prove this, recall (as above, identifying each $f \in B^*$ with its restriction to B) that B^* is a closed subset of $[-1, 1]^B$ which itself is compact, by Tychonoff's Theorem.

7. Completions and Compactifications

Recall that every metric space X has a natural completion $\widehat{X}$. By definition, $\widehat{X}$ is a complete metric space in which X embeds as a dense subspace, i.e. we have an isometric embedding $\iota : X \to \widehat{X}$. Here, 'embedding' means that ι gives a homeomorphism between X and its image $\iota(X) \subseteq \widehat{X}$. 'Isometric embedding' is the stronger condition that ι preserves distance. The pair $(\widehat{X}, \iota)$ enjoys the following universal property: If $f : X \to Y$

is an isometric embedding of X in a complete metric space Y, then there is an isometric embedding $\widehat{f} : \widehat{X} \to Y$ such that the following diagram commutes:

$$\begin{tikzcd} X \arrow[r, "\iota"] \arrow[dr, "f"'] & \widehat{X} \arrow[d, dashed, "\widehat{f}"] \\ & Y \end{tikzcd}$$

and this universal property suffices to define the completion. The uniqueness of the completion follows from this universal property, while its existence is supplied by our construction

$$\widehat{X} = \mathrm{Cauchy}(X)/\!\sim, \quad \iota(x) = (x, x, x, \ldots)$$

where $\mathrm{Cauchy}(X) \subseteq X^\omega$ is the space of all Cauchy sequences in X.

In a similar way, given any topological space X, we would like to define the compactification of X by the appropriate universal property. But what is the 'right' property? For example there are two obvious ways to compactify $\mathbb{R} \simeq (0,1)$: by adding one endpoint or two. The first approach gives the embedding

$$(0,1) \to S^1, \quad t \mapsto e^{2\pi t i}$$

while the second approach gives the embedding

$$(0,1) \to [0,1], \quad t \mapsto t.$$

In general a compactification of X is understood to mean a compact space in which X embeds as a dense subspace. ('Dense' is required in order to achieve minimality of the compactification. For example, the map $(0,1) \to [0,1]^2$, $t \mapsto (0,t)$ embeds the interval $(0,1)$ in a compact space $[0,1]^2$; but we can discard the points (x,y) with $y > 0$ to obtain $[0,1]$ as a compactification of $(0,1)$.) Not every space has a compactification in this sense; and when it exists, it is not necessarily unique, as our examples show.

Given a topological space X, the easiest compactification is the **one-point compactification** $X \cup \{\infty\}$ where we have added just one new point, denoted '∞'. Open sets in $X \cup \{\infty\}$ are of two types:

- open subsets of X remain open in $X \cup \{\infty\}$; and
- subsets of the form $(X \smallsetminus K) \cup \{\infty\}$ where $K \subseteq X$ is closed and compact.

The resulting topological space $X \cup \{\infty\}$ is compact whenever X itself is Hausdorff and locally compact but not compact. (We say that X is **locally compact** if every point has an open neighborhood which is contained in some compact subset.) For example, the one-point compactification of $\mathbb{R}$ is homeomorphic to S^1 (the first example listed above). The second compactification shown above: $(0,1) \to [0,1]$, $t \mapsto t$, is an example of Stone-Čech compactification. This approach takes a little more work to explain in general; but it exists under much more general conditions on the original space X. We will assume only that X is Hausdorff and completely regular.

8. The Stone-Čech Compactification

Let X be a completely regular Hausdorff space (see Appendix A). A **Stone-Čech compactifation** of X is an embedding $\iota : X \to \beta X$ such that for every continuous map from X to a compact Hausdorff space Y, say $f : X \to Y$, there is a *unique* continuous map $\widehat{f} : \beta X \to Y$ such that the following diagram commutes:

$$\begin{tikzcd} X \arrow[r, "\iota"] \arrow[dr, "f"'] & \beta X \arrow[d, dashed, "\widehat{f}"] \\ & Y \end{tikzcd}$$

Here 'embedding' means that ι is a homeomorphism between X and its image $\iota(X) \subseteq \beta X$. We identify X with its image under this embedding, namely $\iota(X) \subseteq \beta X$; and this is necessarily a dense subspace. The uniqueness of the Stone-Čech compactification follows directly from this universal property; it remains for us to prove its existence. We give this proof first in the case that X is discrete. Remarkably, the general proof follows from this special case by taking the appropriate quotient.

Construction of βX when X is discrete

We begin with the case X is a discrete topological space. This special case is in fact the hardest case; using this we will later complete the general case. In the discrete case, we simply define

$$\beta X = \{\mathfrak{U} : \mathfrak{U} \text{ is an ultrafilter on } X\}.$$

Thus points of βX are just ultrafilters on X. We must specify the topology of βX and show that it is compact Hausdorff. A basis for this topology is obtained as follows: for every subset $A \subseteq X$, define

$$[A] = \{\mathfrak{U} \in \beta X : A \in \mathfrak{U}\}.$$

Before proceeding further, recall that in the case of a singleton set $A = \{x\}$, there is a *unique* filter containing $\{x\}$, namely the *principal ultrafilter* which we have denoted

$$\mathfrak{F}_x = \{B \subseteq X : x \in B\}$$

and so in this case, $[\{x\}] = \{\mathfrak{F}_x\}$ is a singleton point in βX.

8.1 Lemma. The map $A \mapsto [A]$, $\mathcal{P}(X) \to \mathcal{P}(\beta X)$ gives an embedding of partially ordered sets under '$\subseteq$'. In particular,

(a) $[\varnothing] = \varnothing$ and $[X] = \beta X$.
(b) If $A \subseteq B \subseteq X$ then $[A] \subseteq [B]$.
(c) $[A \cup B] = [A] \cup [B]$ and $[A \cap B] = [A] \cap [B]$.
(d) $[X \smallsetminus A] = \beta X \smallsetminus [A]$.
(e) The map $\mathcal{P}(X) \to \mathcal{P}(\beta X)$, $A \mapsto [A]$ is one-to-one.

Proof. I will just prove (e), and leave the remaining parts as an exercise. Suppose $A \neq B$ are distinct subsets of X. If there exists $a \in A$ with $a \notin B$, then $B \subseteq X \smallsetminus \{a\}$ so $[B] \subseteq [X \smallsetminus \{a\}]$. However $[A] \supseteq [\{a\}] \neq \varnothing$ is disjoint from $[X \smallsetminus \{a\}]$ by (d). Thus $[A] \neq [B]$. If there exists $b \in B$ with $b \notin A$, the argument is similar. $\square$

8.2 Corollary. The subsets $[A] \subseteq \beta X$ form a basis for a topology on βX. The map $\iota : X \to \beta X$, $x \mapsto \mathfrak{F}_x$ embeds X as a discrete dense subspace of βX. Thus the restriction of ι to $X \to \iota(X)$ is a homeomorphism.

Proof. The fact that the subsets $[A] \subseteq \beta X$ form a basis for a topology on βX follows directly from Lemma 8.1. The map $\iota : X \to \beta X$, $x \mapsto \mathfrak{F}_x$ is clearly one-to-one; and since the singletons $[\{x\}] = \{\mathfrak{F}_x\}$ are open, the image of X under this embedding is discrete. For every nonempty subset $A \subseteq X$, the basic open set $[A] \subseteq \beta X$ satisfies $\iota(x) = \mathfrak{F}_x \in [A]$ whenever $x \in A$; thus the image $\iota(A) \subseteq \beta X$ is dense. The result follows. $\square$

Since $[X \smallsetminus A] = \beta X \smallsetminus [A]$, the sets $[A]$ are clopen; they constitute a family of basic closed sets, as well as a set of basic open sets, in the terminology of Section 1.

8.3 Theorem. For a discrete space X, the space βX as defined above is compact and Hausdorff; moreover it is the Stone-Čech compactification of X.

Proof. Let $\mathfrak{U}, \mathfrak{U}' \in \beta X$ be disticnt points, i.e. distinct ultrafilters on X. For some subset $A \subseteq X$, we have $A \in \mathfrak{U}$ but $A \notin \mathfrak{U}'$, so that $A' \in \mathfrak{U}'$ where $A' = X \smallsetminus A$; then $[A]$ and $[A']$ are disjoint open neighborhoods of the points $\mathfrak{U}, \mathfrak{U}'$ respectively. Thus βX is Hausdorff.

We prove compactness using Lemma 1.2 using the fact that $\{[A] : A \subseteq X\}$ is a family of basic closed sets for βX. Consider an indexed family $\{[A_\alpha]\}_\alpha$ of basic closed sets with the finite intersection property; we must show that $\bigcap_\alpha [A_\alpha] \neq \varnothing$. Since $[A_{\alpha_1} \cap \cdots \cap A_{\alpha_n}] = [A_{\alpha_1}] \cap \cdots \cap [A_{\alpha_n}] \neq \varnothing = [\varnothing]$, by Lemma 8.1(e) we have $A_{\alpha_1} \cap \cdots \cap A_{\alpha_n} \neq \varnothing$. Thus the sets A_α themselves satisfy the finite intersection property; and so they generate a filter on X. Extending this filter to an ultrafilter, we obtain $\mathfrak{U} \in \beta X$ such that $A_\alpha \in \mathfrak{U}$ for all α, i.e. $\mathfrak{U} \in \bigcap_\alpha [A_\alpha]$. This shows that βX is compact.

It remains to be checked that ι satisfies the required universal property. Let $f : X \to Y$ be any map, where Y is an arbitrary compact Hausdorff space. (Since X is discrete, every map defined on X is continuous.) Given $\mathfrak{U} \in \beta X$, the push-forward ultrafilter $f_*(\mathfrak{U})$ on Y must converge to a unique point by Theorem 2.1, and we denote this point $\widehat{f}(\mathfrak{U}) \in Y$. This gives a well-defined function $\widehat{f} : \beta X \to Y$. For all $x \in X$, $\widehat{f}(\iota(x)) = \widehat{f}(\mathfrak{F}_x) = f(x)$ by definition of $\widehat{f}$, since $\mathfrak{F}_x \searrow x$; thus $\widehat{f} \circ \iota = f$.

Since $\iota(X) \subseteq \beta X$ is dense and Y is Hausdorff, it is easy to see that $\widehat{f}$ is the unique map making the defining diagram commute. To see this, let $g : \beta X \to Y$ such that $g \circ \iota = f$, and let $x \in X$.