

B. Ultrafilters exist

Example A2 shows that there exist ultrafilters on all nonempty sets. But the principal ultrafilters exhibited there are, as the alternate terminology "trivial implies", of very little interest. It is therefore natural to ask whether non-principal ultrafilters exist and, if so, on which sets. Such sets must be infinite, because of Proposition A10, (c) \Rightarrow (a). Corollary 2 below asserts that there is no other requirement on such sets.

~~Theorem 1. Let I be any set and let \mathcal{S} be a family of subsets of I . If \mathcal{S} has the finite intersection property then there exists an ultrafilter \mathcal{U} on I such that $\mathcal{S} \subseteq \mathcal{U}$.~~

~~Proof. Apply Zorn's lemma to the ~~family~~ collection, partially ordered by inclusion, of all families \mathcal{S}' of subsets of I such that \mathcal{S}' has the FIP and $\mathcal{S}' \supseteq \mathcal{S}$. Any ~~chain~~ chain in this partially ordered set has an upper bound here~~

Theorem 1. A family \mathcal{S} of subsets of a set I is included in some ultrafilter on I if and only if it has the finite intersection property.

Proof. Since ultrafilters have the FIP, so do all their subfamilies. It remains to prove that every family \mathcal{S} with ^{the} FIP is included in an ultrafilter.

~~Let \mathcal{F}~~ Let \mathcal{Q} be the collection, partially ordered by inclusion of all families S' of subsets of I such that S' has the FIP and $S' \supseteq S$. If Γ is a chain in this partially ordered set, then its union $\bigcup \Gamma$ has the FIP, because any finite subfamily of $\bigcup \Gamma$ is also a finite subfamily of some $S' \in \Gamma$ and therefore has nonempty intersection. Thus, every such chain Γ has an upper bound in \mathcal{Q} , namely S if Γ is empty and $\bigcup \Gamma$ otherwise. By Zorn's lemma, \mathcal{Q} has a maximal element U . ^{Thus, $U \supseteq S$ and} U is a maximal family of subsets of I , with the FIP. By Theorem A13, U is an ultrafilter. \square

In many applications of Theorem 1, S will be a filterbase or a filter. Here is an easy example.

Corollary 2. ~~If I is infinite~~ On every infinite set, there exists a non-principal ultrafilter.

Proof. On an infinite set, the family of cofinite subsets (i.e. the cofinite filter) has the FIP and therefore is included in an ultrafilter. Such an ultrafilter is non-principal by A10. \square

A similar argument, using the last ~~two~~ statement in Digression A18, shows that on every infinite set there exists a uniform ultrafilter.

It is often useful to be able to specify not only that certain sets shall be in an ultrafilter but also that certain sets shall not be in it. This requires no essential extension of Theorem 1, since, by the definition of ultrafilter, to ~~exclude a set from an ultrafilter is the same as~~ keep a set out of an ultrafilter is the same as to put its complement in. Nevertheless, it is convenient to have this version of Theorem 1 stated explicitly.

Corollary 3. Let F and G be families of subsets of I . In order that there exist an ultrafilter U on I that contains all the sets in F but none of the sets in G , it is necessary and sufficient that no ~~finite~~ intersection of finitely many sets in F be included in the union of any finitely many sets in G . If F is a filterbase, then the necessary and sufficient condition reduces to: no set in F is covered by finitely many sets in G .

Proof. An ultrafilter contains all the sets in F and none of the sets in G if and only if it includes

$$S = F \cup \{I - G \mid G \in G\}.$$

By Theorem 1, such an ultrafilter exists if and only if I has the FIP, i.e., if and only if there do not exist $F_1, \dots, F_m \in \mathcal{F}$ and $G_1, \dots, G_m \in \mathcal{G}$ with

$$F_1 \cap \dots \cap F_m \cap (I - G_1) \cap \dots \cap (I - G_m) = \emptyset.$$

This equation is equivalent to

$$F_1 \cap \dots \cap F_m \subseteq G_1 \cup \dots \cup G_m,$$

so the first assertion in the Corollary is established. The second follows because the intersection of finitely many sets from a filterbase includes a set from that filterbase. \square

Corollary 4. Any filter \mathcal{F} on I is the intersection of all the ultrafilters on I that include \mathcal{F} .

Proof. \mathcal{F} is obviously included in this intersection. ~~Further~~ The converse asserts that ~~if $X \in \mathcal{F}$~~ if $X \notin \mathcal{F}$ then there is an ultrafilter \mathcal{U} with $\mathcal{F} \subseteq \mathcal{U}$ but $X \notin \mathcal{U}$. This follows from the preceding corollary, with $\mathcal{G} = \{X\}$, since X , being outside \mathcal{F} , cannot include any subset in \mathcal{F} . \square

Corollary 5. Let \mathcal{F} and \mathcal{H} be families of subsets of I . In order that there exist an ultrafilter \mathcal{U} on I such that $\mathcal{F} \subseteq \mathcal{U} \subseteq \mathcal{H}$ it is necessary and sufficient that

whenever the intersection of finitely many sets from \mathcal{F} is represented as the union of finitely many sets, then at least one of these sets is in \mathcal{H} . If \mathcal{F} is a filterbase, then "the intersection of finitely many sets from \mathcal{F} " can be simplified to "a set from \mathcal{F} ." If \mathcal{H} is closed under taking supersets, then "represented as the union of" can be replaced by "partitioned into".

Proof. The first two assertions are Corollary 3, with $\mathcal{G} = \mathcal{P}(I) - \mathcal{H}$. For the last assertion, observe that if X is represented as the union of Y_1, Y_2, \dots, Y_n , then it is partitioned into $Y_1, Y_2 - Y_1, Y_3 - (Y_1 \cup Y_2), \dots, Y_n - (Y_1 \cup \dots \cup Y_{n-1})$. If \mathcal{H} contains ~~a piece from this~~ the k^{th} piece $Y_k - (Y_1 \cup \dots \cup Y_{k-1})$ of this partition, then it contains Y_k . \square

Zorn's lemma, used in the proof of Theorem 1, is a standard device for efficiently performing certain constructions by transfinite induction. After using transfinite induction once to prove Zorn's lemma (as in []), one avoids a multitude of similar inductions by invoking this lemma instead. Of course, any proof ^{based on} Zorn's lemma can be expressed as a transfinite inductive construction instead, but there is usually no point in such an expansion of the proof. Ultrafilter theory, however, involves many transfinite inductive constructions that are not easily brought into the scope of Zorn's lemma. In preparation for ~~some~~ these constructions, which will be presented in the format of transfinite inductions, it seems useful to consider, as a simple example, ^a proof of ^{the non-trivial half of} Theorem 1 in this format. What follows is such a proof, slightly modified (for example by using filters instead of general families with the FIP) to increase its similarity to later constructions of this sort. It begins with a lemma that will be used in ~~all~~ all these constructions.

Lemma 6. If Φ is a ^{nonempty} family of filters on a set I and if Φ is linearly ordered by inclusion, then the union $\bigcup \Phi$ is also a filter on I .

Proof. In Definition A8 of filters, clauses (i) and (iii) are trivially preserved by arbitrary unions, and nonemptiness is preserved by unions of ^{arbitrary} nonempty families. As for clause (ii) suppose X and Y are in $\bigcup \Phi$. So there are $\mathcal{F}_1, \mathcal{F}_2 \in \Phi$ with $X \in \mathcal{F}_1$ and $Y \in \mathcal{F}_2$. As Φ , linearly ordered, assume, without loss of generality, that $\mathcal{F}_1 \subseteq \mathcal{F}_2$. Then \mathcal{F}_2 contains both X and Y and therefore also $X \cap Y$. So $X \cap Y \in \bigcup \Phi$. \square

Construction 7. Let \mathcal{S} be a family of subsets of I with the FIP. To find an ultrafilter $\mathcal{U} \supseteq \mathcal{S}$ means (by A13) to find a filter $\mathcal{U} \supseteq \mathcal{S}$ such that, for every $X \subseteq I$, the following requirement is satisfied.

Requirement X : Either $X \in \mathcal{U}$ or $I - X \in \mathcal{U}$.

~~The construction of \mathcal{U} will be a transfinite induction, indexed by the requirements, i.e. by $\mathcal{P}(I)$. ~~Each stage of the induction~~ Each stage in the construction will ensure that the corresponding requirement will be satisfied by the final result \mathcal{U} . ~~Each stage~~~~

The construction of \mathcal{U} will be a transfinite induction, in which each stage will ensure that one of these requirements is satisfied by the final result \mathcal{U} .

The construction begins with the filter generated by the given family \mathcal{S} . At the stage corresponding to requirement X , either X or $I-X$ is adjoined to the filter produced by the previous stages. In detail, the construction is as follows.

Fix a transfinite enumeration of the requirements, i.e., a bijection from some ordinal number κ to the set $\mathcal{P}(I)$ of all requirements. (κ can be taken to be the cardinality of $\mathcal{P}(I)$.) The construction will associate to each ordinal $\alpha \leq \kappa$ a filter \mathcal{F}_α in such a way that:

- (a) If $\alpha \leq \beta$ then $\mathcal{F}_\alpha \subseteq \mathcal{F}_\beta$.
- (b) \mathcal{F}_0 is the filter generated by \mathcal{S} .
- (c) If β is a limit ordinal, then $\mathcal{F}_\beta = \bigcup_{\alpha < \beta} \mathcal{F}_\alpha$.
- (d) If $\alpha < \beta$ and the α^{th} requirement is $X \in I$, then either $X \in \mathcal{F}_\beta$ or $I-X \in \mathcal{F}_\beta$.

In the presence of (a), (d) can be replaced by

- (d') If the α^{th} requirement is $X \in I$ then either $X \in \mathcal{F}_{\alpha+1}$ or $I-X \in \mathcal{F}_{\alpha+1}$.

Furthermore, an inductive argument, using ~~(a)~~ (c) at limit stages, shows that (a) can be replaced by

(a') For all $\alpha < \kappa$, $\mathcal{F}_\alpha \subseteq \mathcal{F}_{\alpha+1}$.

The inductive ~~construction~~ definition of \mathcal{F}_α and verification of (a'), (b'), (c), and (d') proceed as follows.

Define \mathcal{F}_0 ~~by~~ by (b). All the other clauses are trivial at this stage.

If β is a limit ordinal, then define \mathcal{F}_β by (c). Again, the other clauses are trivial.

If $\beta = \alpha + 1$, let $Z \subseteq I$ be the α^{th} requirement. Of the two families $\mathcal{F}_\alpha \cup \{Z\}$ and $\mathcal{F}_\alpha \cup \{I - Z\}$, at least one has the FIP, for otherwise ~~(a)~~ \mathcal{F}_α would lack the FIP, as in the proof of A13, (c') \Rightarrow (b). Let $\mathcal{F}_{\alpha+1}$ be the filter generated by $\mathcal{F}_\alpha \cup \{Z\}$ if this has the FIP and the filter generated by $\mathcal{F}_\alpha \cup \{I - Z\}$ otherwise. Then (a') and (d') clearly hold, and the other two clauses are ~~also~~ trivial at this stage.

This completes the construction of the filters \mathcal{F}_α and the verification of (a) through

(d). The special case $\beta = \kappa$ of (d) says that \mathcal{F}_κ satisfies all the requirements $X \in \mathcal{P}(I)$. It is therefore an ultrafilter by A13, and it includes \mathcal{S} by (a) and (b). \square

This construction serves as a prototype for many constructions of ultrafilters with various special properties. Such constructions generally begin by specifying requirements to the effect that the ultrafilter is to contain sets of a certain sort. The set of requirements is enumerated in a transfinite sequence, and a sequence of filters is constructed, by induction, so that (a), (b), and (c) as above hold (often with some specific \mathcal{S} , such as the cofinite filter, in (b)). In addition, an induction hypothesis analogous to (d) is maintained, namely that \mathcal{F}_β contains sets of the sorts demanded by the requirements preceding the β^{th} in the enumeration. ~~Most of the induction proceeds exactly as here; the only alteration occurs in the successor~~

Thus, at successor stages of the induction, it must be verified that \mathcal{F}_α can be extended to a filter $\mathcal{F}_{\alpha+1}$ ~~which~~ satisfying the α^{th} requirement.

Remark 8. In Construction 7, if Z is the α^{th} requirement and if $\mathcal{F}_\alpha \cup \{Z\}$ lacks the FIP, then adjoining $I-Z$ to \mathcal{F}_α in forming $\mathcal{F}_{\alpha+1}$ was redundant, since $I-Z$ already in \mathcal{F}_α . Indeed, the failure of the FIP means that for some $F_1, \dots, F_n \in \mathcal{F}_\alpha$,

$$F_1 \cap \dots \cap F_n \cap Z = \emptyset,$$

so

$$F_1 \cap \dots \cap F_n \subseteq I-Z.$$

Similarly, if $\mathcal{F}_\alpha \cup \{I-Z\}$ lacks the FIP, then Z is in \mathcal{F}_α , so its adjunction in forming $\mathcal{F}_{\alpha+1}$ was redundant. Thus, the only stages β where \mathcal{F}_β contains sets not in any previous \mathcal{F}_α are $\beta=0$ and those successor stages $\alpha+1$ such that, with Z denoting the α^{th} requirement, both $\mathcal{F}_\alpha \cup \{Z\}$ and $\mathcal{F}_\alpha \cup \{I-Z\}$ have the FIP. We call these the essential stages of the construction. Notice that each \mathcal{F}_α is generated by S and the sets $(Z \text{ or } I-Z)$ adjoined at the essential stages prior to α . (Proof by induction on α .) In particular, the ultrafilter $\mathcal{U} = \mathcal{F}_\kappa$ is generated by S together with the sets adjoined at all the essential stages.

Digression 9. Both proofs of Theorem 1 made essential use of the axiom of choice, the one by invoking Zorn's lemma, and the other by well-ordering $P(I)$. It is natural to ask whether ~~the~~ the theorem can be proved without the axiom of choice. It is shown in [] ~~that~~ the usual axioms of set theory, minus the axiom of choice, ~~it~~ cannot. The usual axioms of set theory, minus the axiom of choice, are consistent with "All ultrafilters on the set of natural numbers are principal" [] and even with "All ultrafilters are principal". It is therefore natural to reverse the question and ask whether Theorem 1 implies the axiom of choice. It does not. Even if Theorem 1 is strengthened to refer to ultrafilters in arbitrary Boolean algebras, rather than ~~just~~ ^(as in A19) just power sets, it can hold while the axiom of choice fails.

Variants of these results show that, even if one assumes the axiom of choice so that non-principal ultrafilters exist, it does not follow that any such ultrafilters are definable.

On the other hand, in the absence of the axiom of choice, there can be ultrafilters ~~and~~ even definable ones, quite different from any that are possible when the

axiom of choice holds. For example, there can be infinite sets I on which the cofinite filter is an ultrafilter; this amounts to saying that I cannot be partitioned into two infinite pieces, and it clearly cannot happen when the axiom of choice holds. Such sets I are called amorphous; see [] for more information.

~~In models of set theory where all sets are Lebesgue measurable (which is~~
~~possible if one assumes, instead of the axiom of choice, that all sets of real numbers are~~
 Lebesgue measurable and the axiom of dependent choice holds (which is consistent if inaccessible cardinals are []), then there are ~~no~~ no non-principal ultrafilters on the set ω of natural numbers, by ~~Theorem~~ Theorem K below. But there does exist a non-principal ultrafilter on the set of equivalence classes of ~~subsets of~~ ~~models~~ the equivalence relation " $X \Delta Y$ is finite". Indeed, for ~~any~~ ^{any} subset of K , the preimage under the canonical projection $P(\omega) \rightarrow K$ is measurable with respect ~~to~~ the real numbers models the equivalence relation " $x - y$ is rational". Indeed, for any $X \in K$, the pre-image $f^{-1}(X)$ under the canonical projection of the interval $[0, 1]$ onto K is ~~no~~ Lebesgue measurable by assumption and thus has

measure 0 or 1 by the zero-one law []. The subsets of K for which $f^{-1}(X)$ has measure 1 are easily seen to form a countably complete non-principal ultrafilter on K . See [] for this and other interesting properties of K , e.g. that its cardinality is strictly greater than that of the continuum.

Application 10. Theorem 1 has numerous consequences that do not explicitly involve ultrafilters, ~~among the various~~ including various so-called compactness principles, which assert that, roughly speaking, solutions to certain sorts of ^(infinite) combinatorial problems can be pieced together from solutions to finite subproblems.

The following is a typical such principle.

Rado's selection principle. ~~Let K be a family of functions~~ Let A be an arbitrary set and B a finite set. Let K be a family of functions, each having its domain included in A and its range in B . Assume that ~~each~~ ^{each} finite subset of A is included in the domain of at least one $f \in K$. Then there is a function $g: A \rightarrow B$ such that, for every finite $F \subseteq A$, there is an $f \in K$ whose domain includes F and whose ~~restriction~~ restriction to F coincides with that of g .

Proof. For each finite $F \subseteq A$, let

$$D(F) = \{f \in K \mid \text{The domain of } f \text{ includes } F\}.$$

Since, for any ~~finite~~ ^{finitely many} finite subsets F_1, \dots, F_m of A ,

$$D(F_1) \cap \dots \cap D(F_m) = D(F_1 \cup \dots \cup F_m)$$

is nonempty, by hypothesis, the family $\mathcal{S} = \{D(F) \mid F \text{ finite}, F \subseteq A\}$ has the FIP. By Theorem 1, let \mathcal{U} be an ultrafilter ~~on~~ on K with $\mathcal{S} \subseteq \mathcal{U}$. For each

$a \in A$, the function $f \mapsto f(a)$ is defined on a set in \mathcal{U} , namely $D(\{a\})$, and takes only finitely many values. By A5(c), it is constant ~~on some set~~ ~~with value~~ ~~$g(a)$~~

on some set $X(a) \in \mathcal{U}$; call its constant value there $g(a)$. Then $g: A \rightarrow B$. ~~To show that g is as required, consider~~

~~an arbitrary finite~~ $F = \{a_1, \dots, a_n\} \subseteq A$. ~~The sets~~ $X(a_1), \dots, X(a_n)$

and $D(F)$ are all in \mathcal{U} , so their intersection is nonempty. Let f be a member of this

intersection. Then $f \in K$, the domain of f includes F , and, for each $a_i \in F$,

$f(a_i) = g(a_i)$ because $f \in X(a_i)$. ~~F was arbitrary, this shows the result.~~