Ultrafilters on Countable Sets

Andreas Blass

University of Michigan
Ann Arbor, MI 48109

ablass@umich.edu

Ultrafilters

An *ultrafilter* on a set A is a family $\mathcal{U}$ of subsets of A such that, for all $X, Y \subseteq A$

(1) $X \cap Y \in \mathcal{U} \iff X \in \mathcal{U}$ and $Y \in \mathcal{U}$

(2) $A - X \in \mathcal{U} \iff X \notin \mathcal{U}$.

It follows that

(3) If $X \in \mathcal{U}$ and $X \subseteq Y \subseteq A$ then $Y \in \mathcal{U}$.

(4) $A \in \mathcal{U}$

(5) $\varnothing \notin \mathcal{U}$

(6) $X \cup Y \in \mathcal{U} \iff X \in \mathcal{U}$ or $Y \in \mathcal{U}$

Items (1), (4), and (5) define *(proper) filters*; (3) follows.

Ultrafilters are maximal filters.

Notation: If $\mathcal{Q}$ is any superset-closed family of subsets of A, we write $(\mathcal{Q}x)$ and say "for $\mathcal{Q}$-many x" for the quantifier over A defined by

$$
\begin{aligned}
(\mathcal{Q}x)\,\varphi(x) &\iff \{x \in A : \varphi(x)\} \in \mathcal{Q} \\
&\iff (\exists X \in \mathcal{Q})(\forall x \in X)\,\varphi(x).
\end{aligned}
$$

Other Views of Ultrafilters

Filters on A are the sets $h^{-1}(\{1\})$ for homomorphisms h from $\mathcal{P}(A)$ to Boolean algebras. Ultrafilters are the special case where h maps to $\{0,1\}$.

Ultrafilters on A are the points of the Čech-Stone compactification βA of the discrete space A.

Equivalently, they are operations which, for any compact Hausdorff space C, transform maps $A \to C$ to points of C, respecting continuous maps between compact Hausdorff spaces C.

The nonempty closed sets of βA are the sets

$$\overline{\mathcal{F}} = \{\mathcal{U} \in \beta A : \mathcal{F} \subseteq \mathcal{U}\}$$

for filters $\mathcal{F}$ on A.

Ultrapowers

Ultrafilters on A can be identified with certain operators, ultrapowers, acting on first-order structures and, in particular, on the *complete structure* $\mathfrak{A}$ on A, whose universe is A and whose relations and functions are *all* of the finitary relations and functions on A.

$$\mathcal{U}\text{-prod}\,\mathfrak{X} = \langle X^A/(f \sim g \iff (\mathcal{U}a)\, f(a) = g(a)), \ldots \rangle$$

An elementary extension of $\mathfrak{A}$ is an ultrapower by an ultrafilter on A iff it is generated by a single element.

Any element d in any elementary extension $\mathfrak{D}$ of the complete structure on $\mathfrak{A}$ determines an ultrafilter on A, the *type* of d (in $\mathfrak{D}$), namely $\{X \subseteq A : \mathfrak{D} \models \dot{X}(d)\}$, where $\dot{X}$ is the symbol that denotes X in $\mathfrak{A}$.

The submodel of $\mathfrak{D}$ generated by d is isomorphic to the ultrapower of $\mathfrak{A}$ by this type, via $\dot{f}(d) \mapsto [f]$.

Triviality

Each $a \in A$ gives an ultrafilter

$$\hat{a} = \{X \subseteq A : a \in X\}.$$

Such ultrafilters are called *trivial* or *principal.*

Their quantifiers amount to substitution

$$(\hat{a}x)\,\varphi(x) \iff \varphi(a).$$

As points of βA, they are the elements of A.

$$\hat{a}\text{-prod}\,\mathfrak{X} \cong \mathfrak{X} \quad : \quad [f] \mapsto f(a).$$

As an operation on A-indexed families in compact Hausdorff spaces, $\hat{a}$ picks out the element indexed by a.

An ultrafilter is trivial iff it contains a finite set.

Convention: Whenever necessary, tacitly assume ultrafilters are nontrivial. Also tacitly assume that filters include the cofinite filter

$$\{X \subseteq A : A - X \text{ is finite}\}$$

Rudin-Keisler Ordering

If $f : A \to B$ and $\mathcal{U}$ is an ultrafilter on A, then

$$f(\mathcal{U}) = \{Y \subseteq B : f^{-1}(Y) \in \mathcal{U}\}$$

is an ultrafilter on B.

This is the unique continuous extension of $f : A \to B$ to $\beta A \to \beta B$.

$$(f(\mathcal{U})y)\,\varphi(y) \iff (\mathcal{U}x)\,\varphi(f(x)).$$

There is a canonical elementary embedding

$$f(\mathcal{U})\text{-prod}\,\mathfrak{X} \to \mathcal{U}\text{-prod}\,\mathfrak{X} : [g] \mapsto [g \circ f],$$

and all elementary embeddings between ultrapowers over A of the complete structure on A arise in this way.

Define $\mathcal{V} \leq \mathcal{U}$ to mean that $\mathcal{V} = f(\mathcal{U})$ for some f. This is the *Rudin-Keisler ordering* of ultrafilters. If $\mathcal{V} \leq \mathcal{U}$, we also say that $\mathcal{V}$ is an *image* of $\mathcal{U}$.

Tensor Products

Let $\mathcal{U}$ and $\mathcal{V}$ be ultrafilters on A and B, respectively. Then

$$\{X \subseteq A \times B :$$
$$\{a \in A : \{b \in B : \langle a, b\rangle \in X\} \in \mathcal{V}\} \in \mathcal{A}\}$$

is an ultrafilter on $A \times B$, written $\mathcal{U} \otimes \mathcal{V}$.

$$((\mathcal{U} \otimes \mathcal{V})\langle x, y\rangle)\, \varphi(x, y) \iff (\mathcal{U}x)(\mathcal{V}y)\, \varphi(x, y).$$

$$(\mathcal{U} \otimes \mathcal{V})\text{-prod}\, \mathfrak{X} \cong \mathcal{U}\text{-prod}\, \mathcal{V}\text{-prod}\, \mathfrak{X}.$$

The two projections from $A \times B$ to A and to B send $\mathcal{U} \otimes \mathcal{V}$ to $\mathcal{U}$ and $\mathcal{V}$, respectively.

If $\mathcal{U}$ and $\mathcal{V}$ are nontrivial ultrafilters on ω, then $\mathcal{U} \otimes \mathcal{V}$ contains $\{\langle x, y\rangle \in \omega^2 : x < y\}$, which we may identify with $[\omega]^2$.

In fact, $\mathcal{U} \otimes \mathcal{V}$ contains $\{\langle x, y\rangle \in \omega^2 : f(x) < y\}$ for any $f : \omega \to \omega$.

Selective Ultrafilters

The best ultrafilters on ω are the *selective* ones, characterized by the following equivalent properties.

- Every partition of ω into sets not in $\mathcal{U}$ admits a selector in $\mathcal{U}$.
- Every function on ω becomes one-to-one or constant when restricted to some set in $\mathcal{U}$.
- Every partition of $[\omega]^2$ into two pieces has a homogeneous set in $\mathcal{U}$. [Kunen]
- For any $n, k \in \omega$, every partition of $[\omega]^n$ into k pieces has a homogeneous set in $\mathcal{U}$. [Kunen]
- Every partition of $[\omega]^\omega$ into an analytic piece and a co-analytic piece has a homogeneous set in $\mathcal{U}$. [Mathias]
- Only three ultrafilters on ω^2 project to $\mathcal{U}$ via both projections (namely $\mathcal{U} \otimes \mathcal{U}$, its reflection in the diagonal $\tau(\mathcal{U} \otimes \mathcal{U})$ where $\tau\langle x, y\rangle = \langle y, x\rangle$, and $\Delta(\mathcal{U})$ where $\Delta : \omega \to \omega^2 : x \mapsto \langle x, x\rangle$).

Synonyms for "selective" include "Ramsey" and "Rudin-Keisler minimal."

P-points

P-points on ω are the ultrafilters characterized by the following equivalent properties.

- Every function on ω becomes finite-to-one or constant when restricted to some set in $\mathcal{U}$.
- Every countably many sets $X_n \in \mathcal{U}$ have a *pseudo-intersection* in $\mathcal{U}$, i.e., a set $Y \in \mathcal{U}$ such that $Y - X_n$ is finite for each n.
- In $\beta\omega$, every intersection of countably many neighborhoods of $\mathcal{U}$ is again a neighborhood (not necessarily open) of $\mathcal{U}$.
- For any finite partition of $[\omega]^2$, there exist a set $H \in \mathcal{U}$ and a function $f : \omega \to \omega$ such that all pairs $\{x, y\} \in [H]^2$ with $f(x) \leq y$ lie in the same piece of the partition. [Taylor]
- In $\mathcal{U}$-prod $\mathfrak{N}$, every non-standard elementary submodel is cofinal.
- Every image of $\mathcal{U}$ on a linearly ordered set contains a set whose image has order-type ω or ω^* or 1. [Booth]

Rare and Rapid Ultrafilters

Rare ultrafilters on ω, also called *Q-points*, are those $\mathcal{U}$ such that every finite-to-one function on ω becomes one-to-one when restricted to some set in $\mathcal{U}$.

Equivalently, every element of $\mathcal{U}$-prod $\mathfrak{N}$ that generates a cofinal submodel in fact generates the whole model.

An ultrafilter is selective iff it is both a P-point and rare.

Rapid ultrafilters on ω are those $\mathcal{U}$ such that, for each function $f : \omega \to \omega$, there is some $X \in \mathcal{U}$ whose enumerating function majorizes f, i.e.,

$$|\{x \in X : x < f(n)\}| \leq n \quad \text{for each } n \in \omega.$$

Equivalently, any element of $\mathcal{U}$-prod $\mathfrak{N}$ that generates a cofinal submodel is $\geq$ some element that generates the whole model.

Rare implies rapid.

Existence

The existence of ultrafilters with all these properties is easy to prove assuming CH.

Construction of a selective ultrafilter under CH:

Build an ultrafilter by starting with the cofinite filter and adjoining more sets, in a transfinite induction, to handle all the requirements.

The requirements say that the ultrafilter must contain sets of certain sorts.

There are only $\mathfrak{c} = 2^{\aleph_0}$ requirements, so, by CH, well-order them in a sequence of length $\aleph_1$.

At each stage, the filter built so far is countably generated and therefore has a pseudo-intersection. That makes it easy to handle the next requirement.

The same proof works without CH as long as every filter on ω generated by $< \mathfrak{c}$ sets has a pseudo-intersection (i.e., $\mathfrak{p}$=$\mathfrak{c}$).

In fact, $\mathbf{cov}(B) = \mathfrak{c}$ suffices.

Existence

Under CH, one can also construct ultrafilters with any combination of the properties "selective," "P-point," and "rapid" or negations thereof, except those excluded by the facts that selective implies both P-point and rapid.

In ZFC alone, none of this can be proved. It is consistent that there are no P-points [Shelah], and it is consistent that there are no rapid ultrafilters [Miller].
It is still open whether it is consistent to have neither P-points nor rapid ultrafilters. That would require $\mathfrak{c} \geq \aleph_3$.

Forcing an Ultrafilter

The natural way to adjoin an ultrafilter on ω by forcing uses as conditions the infinite subsets X of ω, ordered by inclusion.

The "meaning" of X is that the generic ultrafilter is to contain X. That suggests that the ordering should be inclusion-mod-finite and that conditions that agree mod finite should be identified.

That produces the separative quotient of the forcing.

It is countably closed, so no new reals are adjoined.

The generic ultrafilter is selective.

The proof of this is "the same" as the inductive step in the CH construction of a selective ultrafilter.

Complete Combinatorics

The generic ultrafilter has no combinatorial properties beyond selectivity.

What does that mean?

Suppose κ is a Mahlo cardinal, and consider the universe obtained by Lévy-collapsing all cardinals $< \kappa$ to ω. (So κ is the new $\aleph_1$.) The resulting *Lévy-Mahlo* model satisfies CH, so it has plenty of selective ultrafilters. It also satisfies the following [Mathias]:
If $\mathcal{U}$ is a selective ultrafilter and if $[\omega]^\omega$ is partitioned into two HOD($\mathbb{R}$) pieces, then there is a homogeneous set in $\mathcal{U}$.
That implies that $\mathcal{U}$ is generic (w.r.t. forcing with $[\omega]^\omega$ as above) over the model HOD($\mathbb{R}$).

We express this by saying that selectivity is *complete combinatorics* for forcing with $[\omega]^\omega$.

Note that this forcing is equivalent to forcing with countably generated filters on ω, ordered by reverse inclusion.

F_σ Filter Forcing

One construction of non-selective P-points, under CH, builds an increasing sequence of approximating filters, which are F_σ subsets of the power set of ω [Mathias, Daguenet].

Without CH, the same combinatorial work applies to generic ultrafilters produced by the forcing where:

- Conditions are F_σ filters.
- Extension means superset.

This forcing is countably closed, so it adds no new reals.
It easily follows that the union of the generic set is an ultrafilter. We call such ultrafilters *F_σ-generic.*

F_σ-generic ultrafilters are

- P-points
- but not rapid
- therefore not selective.

Every image of an F_σ-generic ultrafilter is also F_σ-generic.
So F_σ-generic ultrafilters are P-points with no rapid images.

Mathias Forcing

Plain *Mathias forcing* has

- conditions $\langle s, A\rangle$ where $s \in [\omega]^{<\omega}$ and $A \in [\omega]^\omega$.
- $\langle s', A'\rangle \leq \langle s, A\rangle$ if s is an initial segment of s', $A' \subseteq A$, and $s' - s \subseteq A$.

The "meaning" of $\langle s, A\rangle$ is that the generic set $g \subseteq \omega$ should have s as an initial segment and should, except for s, be $\subseteq A$.

Mathias forcing guided by an ultrafilter $\mathcal{U}$ is the same except that the second components A of conditions $\langle s, A\rangle$ must be in $\mathcal{U}$.

This adds a generic pseudo-intersection of all the (ground-model) sets in $\mathcal{U}$.

Plain Mathias forcing is equivalent to a 2-step iteration: First force with $[\omega]^\omega$ to adjoin a selective ultrafilter $\mathcal{U}$; then do Mathias forcing guided by $\mathcal{U}$.
Since $\mathcal{U}$ is rapid, the enumeration of the generic $g \subseteq \omega$ is a dominating real.

Non-Dominating Mathias Forcing

Are there ultrafilters $\mathcal{U}$ such that Mathias forcing guided by $\mathcal{U}$ does not add a dominating real?

Such *Mathias non-dominating* ultrafilters exist under CH. [Canjar]

They are P-points with no rapid images. [Canjar]

Connections with F_σ-genericity?

Strong P-points

F_σ-generic ultrafilters are *strong P-points.* [Laflamme]

This means that, given countably many closed subsets $\mathcal{C}_n$ of $\mathcal{U}$, there is a partition of ω into intervals $I_n = [i_n, i_{n+1})$ such that, for any choice of $X_n \in \mathcal{C}_n$, for all $n \in \omega$,

$$\bigcup_n (X_n \cap I_n) \in \mathcal{U}.$$

Every strong P-point is a P-point with no rapid images. [Laflamme]

Every Mathias non-dominating ultrafilter is a strong P-point. [Laflamme]

Every F_σ-generic ultrafilter is Mathias non-dominating. [Canjar]

Mathias Non-Dominating

An ultrafilter $\mathcal{U}$ is Mathias non-dominating iff it satisfies the *Hrušák-Minami* condition (HM), defined as follows.

For any filter $\mathcal{F}$ on a set A, let

$$\begin{aligned}\mathcal{F}^+ &= \{X \subseteq A : A - X \notin \mathcal{F}\} \\ &= \{X \subseteq A : (\forall Y \in \mathcal{F})\, X \cap Y \neq \varnothing\}.\end{aligned}$$

For an ultrafilter $\mathcal{U}$ on ω, let $\mathcal{U}^{<\omega}$ be the filter on $[\omega]^{<\omega} - \{\varnothing\}$ generated by the sets $[X]^{<\omega} - \{\varnothing\}$ for $X \in \mathcal{U}$.

Condition (HM) says that any countable decreasing sequence of sets from $(\mathcal{U}^{<\omega})^+$ has a pseudo-intersection in $(\mathcal{U}^{<\omega})^+$. ($\mathcal{U}^{<\omega}$ is a *P^+-filter*.)

$\mathcal{U}$ satisfies (HM) iff $\mathcal{U}$ is a strong P-point.

Under CH, there are P-points that have no rapid images but are not strong P-points. [Hrušák, Verner]

Question: Might (HM) be complete combinatorics for F_σ-genericity?

Laflamme Picture of F_σ Filters

Let $\mathcal{Q}$ be an upward-closed family of subsets of A. The *intersection number* of $\mathcal{Q}$ is the smallest $n \in \omega$ such that some n members of $\mathcal{Q}$ have empty intersection.

A *Laflamme sequence* is a sequence $\langle I_n, \mathcal{Q}_n \rangle_{n\in\omega}$ of pairs such that:

- the $I_n = [i_n, i_{n+1})$ form a partition of ω into intervals,
- each $\mathcal{Q}_n$ is an upward-closed family of subsets of I_n, and
- the intersection numbers of the $\mathcal{Q}_n$'s tend to ∞ with n.

To a Laflamme sequence, associate the filter

$$\mathcal{L}(\langle I_n, \mathcal{Q}_n \rangle) = \{X \subseteq \omega : (\forall^\infty n)\, X \cap I_n \in \mathcal{Q}_n\}.$$

This is an F_σ filter on ω.

Every F_σ filter on ω is included in one of this form.

So forcing with F_σ filters is equivalent to forcing with Laflamme sequences.

Algebra of Ultrafilters

Any binary operation $*: A \times A \to A$ on a set A extends to βA by

$$\mathcal{U} * \mathcal{V} = *(\mathcal{U} \otimes \mathcal{V}).$$

Equivalently,

$$((\mathcal{U} * \mathcal{V})z)\, \varphi(z) \iff (\mathcal{U}x)(\mathcal{V}y)\varphi(x * y).$$

Equivalently, first extend

$$*: A \times A \to A \subseteq \beta A$$

to

$$*: A \times \beta A \to \beta A$$

by continuity for fixed first argument in A. Then extend to

$$*: \beta A \times \beta A \to \beta A$$

by continuity for fixed second argument in βA.

If $*$ is associative on A, then also on βA. Not so for commutativity.

The operation $*$ on βA is a continuous function of the left argument, with the right argument fixed, but not vice versa.

Idempotents

Let $\langle S, * \rangle$ be a compact semi-topological semigroup: S is a compact Hausdorff space, $*$ is an associative operation on S, and $*$ is continuous in the left argument.

Then $\langle S, * \rangle$ has an idempotent element, i.e., s with $s * s = s$. [Ellis, Numakura]

In particular, for any semigroup $\langle A, * \rangle$, there are idempotent ultrafilters in βA.

This gives a relatively short proof, due to Galvin and Glazer, of Hindman's Theorem: Let the set $\mathbb{F} = [\omega]^{<\omega} - \varnothing$ be partitioned into finitely many pieces. Then there is an infinite set $H \subseteq \mathbb{F}$, such that all finite nonempty unions of members of H are in the same piece of the given partition. Furthermore, the members of H can be taken to be pairwise disjoint.

Proof of Hindman's Theorem

The associative binary operation $\cup$ on $\mathbb{F}$ extends to $\dot\cup$ on $\beta\mathbb{F}$.

Consider those $\mathcal{U} \in \beta\mathbb{F}$ that concentrate on sets that start late, i.e.,

$$(\forall n \in \omega)(\mathcal{U}x)\ \min(x) > n.$$

These $\mathcal{U}$ form a compact subsemigroup of $\langle \beta\mathbb{F}, \dot\cup \rangle$, so there is an idempotent among them. Let $\mathcal{U}$ be such an idempotent.

Given a partition of $\mathbb{F}$ into finitely many pieces, let X be the piece in $\mathcal{U}$. We'll choose elements $x_0, x_1, \dots \in X$ such that

- $\max(x_i) < \min(x_{i+1})$ for every i, and
- every finite union of x_i's is in X.

Because $\mathcal{U}$ is idempotent, whenever we have $(\mathcal{U}x)\,\varphi(x)$, we also have $(\mathcal{U}x)(\mathcal{U}y)\,\varphi(x \cup y)$.
In particular, we have $(\mathcal{U}x)\,x \in X$ and $(\mathcal{U}x)(\mathcal{U}y)\,x \cup y \in X$.

Choose $x_0 \in X$ such that $(\mathcal{U}x)\,x_0 \cup x \in X$.
From this and $(\mathcal{U}x)\,x \in X$ get

$$(\mathcal{U}x)(\mathcal{U}y)\,\left[x \cup y \in X \wedge x_0 \cup x \cup y \in X\right].$$

Also, $(\mathcal{U}x)\ \min(x) > \max(x_0)$.

Proof of Hindman's Theorem

Choose $x_1 \in X$ such that

- $\min(x_1) > \max(x_0)$,
- $x_0 \cup x_1 \in X$, and
- $(\mathcal{U}x)\left[x_1 \cup x \in X \wedge x_0 \cup x_1 \cup x \in X\right]$.

In general, choose x_n so that

- $\min(x_n) > \min(x_{n-1})$,
- for all $t \subseteq \{0, \dots, n-1\}$, $\bigcup_{i \in t} x_i \cup x_n \in X$, and
- for all such t, $(\mathcal{U}x) \bigcup_{i \in t} x_i \cup x_n \cup x \in X$.

Such an x_n exists; in fact there are $\mathcal{U}$-many choices for x_n. After choosing x_n, use idempotence of $\mathcal{U}$ to get, for all t as above,

$$(\mathcal{U}x)(\mathcal{U}y) \bigcup_{i \in t} x_i \cup x_n \cup x \cup y \in X.$$

That, plus its analogs from previous steps, plus "starting late," plus $X \in \mathcal{U}$ provide what is needed for the next step, choosing x_{n+1}.

After ω steps, the chosen x_i's constitute the required homogeneous set.

Union Ultrafilters

Notation: For a set or sequence s of members of $\mathbb{F}$, let $FU(s)$ be the collection of all nonempty finite unions of members of s.

A *union ultrafilter* is an ultrafilter on $\mathbb{F}$ generated by sets of the form $FU(s)$ where s is an infinite sequence of pairwise disjoint members of $\mathbb{F}$.

If the elements s_n of the s can be taken to be ordered in the sense that $\max(s_n) < \min(s_{n+1})$ then we call $\mathcal{U}$ an *ordered union ultrafilter*.

CH implies (using Hindman's theorem) the existence of ordered union ultrafilters.

Union ultrafilters are idempotent with respect to $\dot\cup$.

Consider the two maps $\min, \max : \mathbb{F} \to \omega$. Like any maps, they can be applied to ultrafilters on their domain.

If $\mathcal{U}$ is a union ultrafilter, then $\min(\mathcal{U})$ and $\max(\mathcal{U})$ are P-points [Hindman, A.B.] with no common image [A.B.].

In particular, ZFC doesn't prove the existence of union ultrafilters.

Forcing Union Ultrafilters

The most natural forcing to add an ordered union ultrafilter has as conditions infinite sequences $s = (s_n)_{n\in\omega}$ from $\mathbb{F}$ with $\max(s_n) < \min(s_{n+1})$ for all n.

The "meaning" of s is that $FU(s)$ is in the generic ultrafilter.

s' is an extension of s if $FU(s') \subseteq FU(s)$.

Equivalently, each s'_n is a finite union of s_i's.

We call s' a *condensation* of s.

An extension of s is obtained by deleting some of its terms s_i and merging some of those that survive.

The separative quotient identifies s with s' if some final segments are equal.

This is countably closed, so the forcing adds no new reals.

If G is a generic set of conditions, then $\{FU(s) : s \in G\}$ is a basis for an ordered union ultrafilter on $\mathbb{F}$.

That it is an ultrafilter is a consequence of Hindman's theorem. The combinatorics is the same as in the CH construction of an ordered union ultrafilter.

Stable Ordered Union Ultrafilters

The ordered union ultrafilters obtained by the CH construction or by the forcing described above have the additional property of *stability*, defined as follows:

Given any countably many sequences $s^{(n)}$ with $FU(s^{(n)}) \in \mathcal{U}$, there is an s such that $FU(s) \in \mathcal{U}$ and, for each n, a tail of s is a condensation of $s^{(n)}$.

Note the analogy between this definition and the pseudo-intersection characterization of P-points.

Yet stable ordered union ultrafilters have stronger properties, analogous to those of selective ultrafilters.

Stability and Partitions

Let $\mathcal{U}$ be a stable ordered union ultrafilter. Every function f defined on $\mathbb{F}$ has one of the following five forms on some set in $\mathcal{U}$.

- a constant
- $g \circ \min$ with g one-to-one on ω
- $g \circ \max$ with g one-to-one on ω
- $g \circ \langle \min, \max \rangle$ with g one-to-one on ω^2
- a one-to-one function.

For every partition of the ordered n-sequences from $\mathbb{F}$ into k pieces (with $n, k \in \omega$), there is $H \in \mathcal{U}$ whose ordered n-sequences are all in the same piece.

For every partition of the ordered ω-sequences from $\mathbb{F}$ into an analytic piece and a co-analytic piece, there is $H \in \mathcal{U}$ whose ordered ω-sequences are all in the same piece.

"Stable ordered union" is complete combinatorics for the ordered union forcing described above.

Other Ordered Union Ultrafilters

For ordered union ultrafilters, a definition analogous to P-point (stability) produced results analogous to selectivity.

What is the "real" analog of P-point?

Recall that F_σ-filter forcing produced non-selective P-points.

Is there an analog for ordered union ultrafilters, or more generally for idempotent ultrafilters?

Krautzberger initiated a study of forcing by F_σ idempotent filters.

He showed that, if $\mathcal{U}$ is generic for this forcing, then $\min(\mathcal{U})$ and $\max(\mathcal{U})$ are F_σ-generic, and in fact mutually generic, by product forcing.

Minimal Idempotents

Partially order the idempotents in the semi-group $\langle \beta A, * \rangle$ by

$$\mathcal{U} \preceq \mathcal{V} \iff \mathcal{U} * \mathcal{V} = \mathcal{V} * \mathcal{U} = \mathcal{U}.$$

The following are equivalent, for any idempotent $\mathcal{U}$.

- $\mathcal{U}$ is minimal with respect to $\preceq$.
- $\mathcal{U}$ belongs to some minimal (closed) left ideal of βA.
- $\mathcal{U}$ belongs to all 2-sided ideals of βA.

Such $\mathcal{U}$ are called *minimal idempotents.*
They exist.
In fact, every idempotent is $\succeq$ a minimal one.
Sets that belong to a minimal idempotent ultrafilter, called *central* sets, have strong combinatorial properties.

Hales-Jewett Theorem

Let W be the set of words on a finite alphabet Σ.

Let A be the set of words on $\Sigma \cup \{v\}$ where $v \notin \Sigma$.

Let $V = A - W$.

For $x \in A$ and $\sigma \in \Sigma$, let $x(\sigma) \in W$ be the result of replacing every v in x by σ.

Hales-Jewett Theorem: If W is partitioned into finitely many pieces, then there is some $x \in V$ such that all its instances $x(\sigma)$ are in the same piece of the partition.

In fact, the things substituted for v can be any pre-specified finite subset of A, not necessarily just Σ.

If $\mathcal{U}$ is a minimal idempotent in $\langle \beta W, \frown \rangle$, then the partition piece that is in $\mathcal{U}$ works in the conclusion of the theorem.

Proof of Hales-Jewett

Let $\mathcal{U}$ be a minimal idempotent in βW and P the piece of the partition in $\mathcal{U}$.
Let $\mathcal{V} \preceq \mathcal{U}$ be a minimal idempotent in βA.
Because V is a 2-sided ideal in A, βV is a 2-sided ideal in βA.
So $\mathcal{V} \in \beta V$.
For each $\sigma \in \Sigma$,

$$\hat{\sigma} : A \to W : x \mapsto x(\sigma)$$

is a homomorphism.
It induces a homomorphism $\hat{\sigma} : \beta A \to \beta W$.
Since $\mathcal{V} \preceq \mathcal{U}$, we get that $\hat{\sigma}(\mathcal{V})$ is an idempotent $\preceq \hat{\sigma}(\mathcal{U}) = \mathcal{U}$ in βW.
By minimality, $\hat{\sigma}(\mathcal{V}) = \mathcal{U}$.
So $\hat{\sigma}^{-1}(P) \in \mathcal{V}$.
Pick $x \in \bigcap_{\sigma\in\Sigma} \hat{\sigma}^{-1}(P)$.
This x clearly works.