ULTRAFILTERS AND CARDINAL CHARACTERISTICS OF THE CONTINUUM

ANDREAS BLASS AND NICHOLAS RUPPRECHT

This document is based on lectures given by Andreas Blass at the Tenth Appalachian Set Theory Workshop on January 30, 2010. Nicholas Rupprecht took notes during the workshop and then compiled them into the first version of this document. The final version was prepared by Andreas Blass in the light of useful comments from the referee.

1. Ultrafilters

Convention 1. As is customary in set theory, we identify each natural number with the set of its predecessors; in particular, $2 = \{0, 1\}$. We also follow the set-theoretic convention that the set of all natural numbers is denoted by ω. For any sets X and Y, we write X^Y for the set of all functions from Y into X. This convention creates an ambiguity when the sets X and Y are numbers, but the context will always resolve the ambiguity. The cardinality of a set X is denoted by $|X|$. The collection of all subsets of a set X, the power set of X, is denoted by $\mathcal{P}(X)$ and is identified, via characteristic functions, with 2^X. The collection of subsets of X that have a specified cardinality k is denoted by $[X]^k$.

We write $\mathfrak{c}$ for the cardinality $2^{\aleph_0}$ of the continuum. The continuum hypothesis, $\mathfrak{c} = \aleph_1$, is abbreviated as CH.

When we consider X^Y as a topological space, the intended topology is the product topology obtained from the discrete topology on X. We topologize $\mathcal{P}(Y)$ by its identification with 2^Y, and we topologize $[\omega]^\omega$ as a subspace of $\mathcal{P}(\omega)$. Thus, two elements of $[\omega]^\omega$ are close to each other if they have a long initial segment in common.

Definition 2. An *ultrafilter* on a set X is a family $\mathcal{U}$ of subsets of X such that

(1) $\varnothing \notin \mathcal{U}$;
(2) $X \in \mathcal{U}$;

Partially supported by NSF grant DMS-0653696. The final revision of this paper was prepared while Blass was a visiting fellow in the "Mathematical, Foundational and Computational Aspects of the Higher Infinite" program at the Isaac Newton Institute for Mathematical Sciences.

(3) if $A \subseteq B \subseteq X$ and $A \in \mathcal{U}$ then $B \in \mathcal{U}$;
(4) if $A, B \in \mathcal{U}$, then $A \cap B \in \mathcal{U}$;
(5) for all $A \subseteq X$, either $A \in \mathcal{U}$ or $X - A \in \mathcal{U}$;
(6) if $A \cup B \in \mathcal{U}$, then $A \in \mathcal{U}$ or $B \in \mathcal{U}$.

Remark 3. Some of these six conditions are redundant. The first four alone define the notion of *filter* on X. In (4) and (6), the converse implications are also true by (3). In (5), only one of the two alternatives can occur, because of (1) and (4). Thus, $X - A \in \mathcal{U}$ if and only if $A \notin \mathcal{U}$. We can summarize these facts as: membership in $\mathcal{U}$ respects Boolean combinations (propositional connectives).

Another way to view an ultrafilter $\mathcal{U}$ on X is as a map from the set 2^X to 2, or more generally $M^X \to M$ for any finite set M, namely the map sending any $f : X \to M$ to the unique $m \in M$ for which $f^{-1}(\{m\}) \in \mathcal{U}$. The fact that $\mathcal{U}$ respects Boolean combinations generalizes: As a function $M^X \to M$, it commutes with any k-ary operation $M^k \to M$ and its canonical extension to $(M^X)^k \to M^X$. Conversely, if $\mathcal{U}$ commutes with unary and binary operations on 2, then it is an ultrafilter, because all k-ary operations $2^k \to 2$ for higher k can be expressed in terms of unary and binary operations. An interesting result of Lawvere ([25], see [10, Appendix] for a proof) says that, if we instead consider maps from 3^X to 3, then any map which commutes with just the unary operations comes from an ultrafilter.

A third way to view ultrafilters is as new quantifiers by defining

$$(\mathcal{U}x)\,\varphi(x) \iff \{x \in X : \varphi(x)\} \in \mathcal{U}.$$

The quantifier $(\mathcal{U}x)$ can be read as "for almost all x (with respect to $\mathcal{U}$)" or "for $\mathcal{U}$-most x". Ultrafilter quantifiers (and only those) respect all propositional connectives.

A fourth view of ultrafilters is as uniform ways to choose limits of sequences. Given an X-indexed family $\{p_x\}_{x \in X}$ of points in a compact Hausdorff space C and given an ultrafilter $\mathcal{U}$ on X, there is a unique *limit* of the family along $\mathcal{U}$, denoted by $\mathcal{U}\text{-}\lim_x p_x$, such that every neighborhood of it contains p_x for almost all x. In this way, $\mathcal{U}$ can be viewed as an operation that sends every X-indexed family in any compact Hausdorff space to a point in that space and that commutes with continuous functions. Conversely, every such operation comes from a unique ultrafilter in this way.

A fifth view is that an ultrafilter on X is a point in the Stone-Čech compactification βX of the discrete space X.

Definition 4. An ultrafilter $\mathcal{U}$ on X is said to be *principal* or *trivial* if it is $\{A \subseteq X : q \in A\}$ for some $q \in X$.

The terminology “principal” is imported from ring theory; an ultrafilter is principal if and only if it is generated by a single set, namely $\{q\}$ where q is as in the definition. Also, an ultrafilter is principal if and only if it contains a finite set, and so an ultrafilter is non-principal if and only if it contains all cofinite sets. In terms of the various ways to view ultrafilters, as described above, the principal ultrafilter generated by $\{q\}$ is

- the “evaluation” function $M^X \to M$ that sends each $f \in M^X$ to $f(q)$,
- the “substitution” quantifier, where $(\mathcal{U}x)\,\varphi(x)$ means $\varphi(q)$,
- the operation taking any X-indexed family $\{p_x\}_{x\in X}$ to p_q, and
- the point $q \in X$ regarded as being in βX since $X \subseteq \beta X$.

Ultrafilters on X are exactly the maximal (with respect to $\subseteq$) filters on X. It easily follows, via Zorn's Lemma, that every filter on X is a subset of some ultrafilter on X. In fact, Zorn's Lemma easily gives the following slightly stronger result.

Proposition 5. *Every filter on a set X is the intersection of all the ultrafilters on X that extend it.*

Like any proof using Zorn's Lemma, the argument for existence of ultrafilters can be recast as a transfinite induction, which may shed more light on what is going on in the construction. Well-order the power set $\mathcal{P}(X)$ (the set of all subsets of X) and proceed by transfinite recursion along this well-ordering, starting with a given filter. At the step labeled by a particular set $A \subseteq X$, decide (if it hasn't already been decided) whether to put A or $X - A$ into the ultrafilter under construction. At the end of the inductive process, we have an ultrafilter.

Since the construction appears to involve $2^{|X|}$ choices, it is reasonable to expect the following result of Pospíšil [30].

Proposition 6. *The number of ultrafilters on an infinite set X is $2^{2^{|X|}}$.*

Proof. We invoke a theorem of Hausdorff that there is a family $\mathcal{F}$ of $2^{|X|}$ subsets of X such that, for any disjoint finite subfamilies $\mathcal{A}$ and $\mathcal{B}$ of $\mathcal{F}$, the intersection of the sets in $\mathcal{A}$ and the complements of the sets in $\mathcal{B}$ is nonempty.

Hausdorff's construction begins by saying it is enough to do this with X replaced by

$$X' := \{(P, Q) : P \subseteq X \text{ is finite}, Q \subseteq \mathcal{P}(P)\},$$

since $|X| = |X'|$. Let

$$\mathcal{F} = \{\{(P, Q) \in X' : Y \cap P \in Q\} : Y \subseteq X\}.$$

Then, given disjoint finite subfamilies, say $\{\{(P,Q) : Y \cap P \in Q\} : Y \in \mathcal{I}_+\}$ and $\{\{(P,Q) : Y \cap P \in Q\} : Y \in \mathcal{I}_-\}$, we must find $(P,Q) \in X'$ such that $Y \cap P \in Q$ for all $Y \in \mathcal{I}_+$ and $Y \cap P \notin Q$ for all $Y \in \mathcal{I}_-$. Simply choose P so that all the (finitely many) relevant intersections $Y \cap P$ are distinct, and then choose Q to consist of those intersections where $Y \in \mathcal{I}_+$.

Now, given a family $\mathcal{F}$ (of subsets of X) as in Hausdorff's theorem, observe that, for each $\mathcal{G} \subseteq \mathcal{F}$, there is a filter containing the sets from $\mathcal{G}$ and the complements of the sets from $\mathcal{F} - \mathcal{G}$. Each of these filters can be extended to an ultrafilter; all these ultrafilters, for different $\mathcal{G}$'s, are distinct, and there are as many of them as there are $\mathcal{G}$'s, namely $2^{2^{|X|}}$.

This proves that there are at least $2^{2^{|X|}}$ ultrafilters on X; there cannot be more because an ultrafilter is a subset of $\mathcal{P}(X)$, which has only $2^{2^{|X|}}$ subsets altogether. $\square$

2. Some Partition Theorems

As an application of the quantifier view of ultrafilters we prove (without much work) a weak version of Ramsey's theorem. Afterward, with more work, we prove a much stronger result of Nash-Williams, which implies the full infinitary Ramsey theorem and more. For $\alpha \leq \omega$ and $c, k \in \omega$, let $\omega \to (\alpha)^k_c$ be the statement that, if $[\omega]^k$ is partitioned into c pieces, then there is an $H \subseteq \omega$ such that $|H| = \alpha$ and $[H]^k \subseteq$ one piece.

Theorem 7. *For any natural numbers n, k, and c, $\omega \to (n)^k_c$.*

Proof. As a preliminary step, fix a nonprincipal ultrafilter $\mathcal{U}$ on ω and notice the implications

$$(\forall x > q)\, \varphi(x) \implies (\mathcal{U}x)\, \varphi(x) \implies (\exists x > q)\, \varphi(x)$$

for every $q \in \omega$.

Now to prove the theorem, view k-element subsets of ω as increasing k-tuples. Denote the pieces of the partition by $C_1, \dots, C_c$. Then

$$\forall x_1 \forall x_2 > x_1 \cdots \forall x_k > x_{k-1} \bigvee_{i=1}^{c} \{x_1, \dots, x_k\} \in C_i$$

By the implications noted above, we can replace each $\forall$ with $\mathcal{U}$, and the statement is still true. Now push the disjunction over i to the outside, using the fact that ultrafilter quantifiers respect propositional connectives. So

$$\bigvee_{i=1}^{c} (\mathcal{U}x_1)(\mathcal{U}x_2)\cdots(\mathcal{U}x_k)\{x_1,\dots,x_k\} \in C_i.$$

Fix an i such that the i^{th} disjunct is true. Rewrite this disjunct by renaming variables and introducing dummy variables as follows. For any subset $\{r_1 < \ldots < r_k\}$ of size k in $\{1,\dots,n\}$, let x_l be replaced with y_{r_l}. Let the remaining y_j $(1 \le j \le n)$ be dummy variables, but include quantifiers $(\mathcal{U}y_j)$ over these variables along with the "important" variables y_{r_i}. There are as many such formulas as there are k-element subsets of $\{1,\dots,n\}$; they all say the same thing as the i^{th} disjunct above, but they say it using different choices of the active variables from among $y_1,\dots,y_n$. Form the (highly redundant) conjunction of all of these,

$$\bigwedge_{\{r_1<\cdots<r_k\}\subseteq\{1,\dots,n\}} (\mathcal{U}y_1)\cdots(\mathcal{U}y_n)\{y_{r_1},\dots,y_{r_k}\} \in C_i.$$

Using again that ultrafilter quantifiers respect propositional connectives, pull all the quantifiers out of the conjunction. Then use the implication from $\mathcal{U}$ to $\exists$, noted at the start of this proof, to replace all the ultrafilter quantifiers with existential quantifiers in the form

$$\exists y_1\, \exists y_2 > y_1 \cdots \exists y_n > y_{n-1} \bigwedge_{\{r_1<\cdots<r_k\}\subseteq\{1,\dots,n\}} \{y_{r_1},\dots,y_{r_k}\} \in C_i.$$

This says that there are $y_1 < y_2 < \cdots < y_n$ such that all k-element subsets of $\{y_1, y_2, \dots, y_n\}$ lie in the same piece C_i of the given partition. □

The preceding proof relied mainly on the formal properties of ultrafilter quantifiers, in particular their respecting propositional connectives. With more hands-on work, ultrafilters can yield far stronger partition theorems, such as the following version of a result of Nash-Williams. (Recall that $[\omega]^\omega$ is topologized as a subspace of $\mathcal{P}(\omega)$, which is in turn topologized by identifying it with the product 2^ω of discrete spaces.)

Theorem 8. *If $[\omega]^\omega$ is partitioned into an open piece and a closed piece, then there is an infinite $H \subseteq \omega$ such that $[H]^\omega$ is in one piece.*

Proof. Instead of thinking of infinite subsets of ω, we think of infinite increasing sequences. So $[\omega]^\omega$ is identified as the set of paths through $\omega^{\uparrow<\omega}$, the tree of finite increasing sequences of natural numbers. Following the definition of the topology on $[\omega]^\omega$, we find that there is a set M of nodes in the tree $\omega^{\uparrow<\omega}$ such that the open piece of our partition consists of the paths that go through a node in M.

Fix a nonprincipal ultrafilter $\mathcal{U}$ on ω.

Mark nodes with ordinals, by the following inductive procedure. Mark 0 on those nodes which have initial segments in M. For $\alpha > 0$, mark α on a node s if and only if s has not been marked with an ordinal smaller than α and $(\mathcal{U}n)\,(s^\frown\langle n\rangle$ has been marked with some $\beta < \alpha)$ (where β can depend on n).

Case 1: The empty sequence gets marked with some ordinal α.

Pick, in succession, numbers $h_1 < h_2 < h_3 \cdots$ such that when you choose h_n, for every $F \subseteq \{h_1, \dots, h_{n-1}\}$, $F \cup \{h_n\}$ has a lower mark than F (unless F had mark 0). We can always do this, since if F had a nonzero mark, $\mathcal{U}$-most h_n will work, and there are only finitely many F's. Any infinite subset of the h_i is in the open piece, because, along the path in $\omega^{\uparrow<\omega}$ given by this infinite subset, the marks keep decreasing until they hit 0.

Case 2: The empty sequence is not marked.

Note that if a node is unmarked, $\mathcal{U}$-most of its immediate successors are unmarked. So we may pick $h_1 < h_2 < h_3 < \cdots$ inductively such that all finite subsequences are unmarked (similarly to the previous case). For any infinite subset of $\{h_1, h_2, \dots\}$, all initial segments are unmarked. In particular, they are not marked with 0, so the set is in the closed piece. $\square$

The preceding theorem immediately implies the following corollary, the infinite Ramsey Theorem.

Corollary 9. *For any $k, c \in \omega$, we have $\omega \to (\omega)^k_c$.*

Proof. An easy induction reduces the general case to the case $c = 2$. So suppose $[\omega]^k$ is partitioned into C_0 and C_1. Partition $[\omega]^\omega$ into D_0 and D_1, where $H \in D_i$ if and only if the set of the first k elements of H is in C_i. The pieces D_0 and D_1 are clopen, so there is an infinite H such that, for some i, $[H]^\omega \subseteq D_i$, and hence $[H]^k \subseteq C_i$. $\square$

3. Ultrafilter Constructions

In this section, we describe some ways of generating new ultrafilters from given ones.

Definition 10. Let $\mathcal{U}$ be a family of subsets of X and let $f : X \to Y$. Then $f(\mathcal{U})$ is defined to be the family

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

It is easy to check that if $\mathcal{U}$ is a filter, then so is $f(\mathcal{U})$, and if $\mathcal{U}$ is an ultrafilter, then so is $f(\mathcal{U})$. For the time being, we are concerned only with the ultrafilter case, but the more general situation will arise later.

In terms of quantifiers, we have

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

The unique continuous extension of $f : X \to Y$ to a map of Stone-Čech compactifications $\beta X \to \beta Y$ sends $\mathcal{U}$ to $f(\mathcal{U})$. In terms of operations on families in compact Hausdorff spaces, $f(\mathcal{U})$ acts like $\mathcal{U}$ after a re-indexing along f; that is,

$$f(\mathcal{U})\text{-}\lim_y p_y = \mathcal{U}\text{-}\lim_x p_{f(x)}.$$

This notion of the image of an ultrafilter under a function leads to the Rudin-Keisler ordering,[1] defined as follows.

Definition 11. An ultrafilter $\mathcal{U}$ on X is *Rudin-Keisler above* an ultrafilter $\mathcal{V}$ on Y, written $\mathcal{U} \geq_{RK} \mathcal{V}$, if $\mathcal{V} = f(\mathcal{U})$ for some $f : X \to Y$.

This ordering is reflexive and transitive, but not antisymmetric (so not a partial order). We do, however, have the following result.

Proposition 12. *If $\mathcal{V} \leq_{RK} \mathcal{U} \leq_{RK} \mathcal{V}$, then there is an f such that $f(\mathcal{U}) = \mathcal{V}$ and f is one-to-one on a set in $\mathcal{U}$.*

Proof. We use the following general lemma, apparently first published as a problem [21].

Lemma 13. *If $f : X \to X$, then X can be decomposed as a disjoint union $X = A_0 \cup A_1 \cup A_2 \cup A_3$ such that $f \restriction A_0$ is the identity, and for $i = 1, 2, 3$, $f(A_i) \cap A_i = \varnothing$.*

Corollary 14. *If $\mathcal{U}$ is an ultrafilter on X and if $f : X \to X$ satsifies $f(\mathcal{U}) = \mathcal{U}$, then there is a set in $\mathcal{U}$ on which f is the identity function.*

Proof. Let A_0, A_1, A_2, A_3 be as in the lemma. Being an ultrafilter, $\mathcal{U}$ must contain one of the A_i; fix that value of i. But then, as $f(\mathcal{U}) = \mathcal{U}$, we also have $f^{-1}(A_i) \in \mathcal{U}$. In particular, A_i and $f^{-1}(A_i)$ cannot be disjoint. So $i = 0$, $A_0 \in \mathcal{U}$, and f is the identity on A_0, as required. $\square$

Now to prove the proposition, let f and g witness $\mathcal{V} \leq_{RK} \mathcal{U}$ and $\mathcal{U} \leq_{RK} \mathcal{V}$, respectively. Apply Corollary 14 to $g \circ f$, which maps $\mathcal{U}$ to itself. We get a set $A \in \mathcal{U}$ on which $g \circ f$ is the identity and therefore f is one-to-one. $\square$

[1]There is a more general ordering, the Katětov ordering, defined on filters as follows. Put $\mathcal{F}$ above $\mathcal{G}$ in the Katětov order if $f(\mathcal{F}) \supseteq \mathcal{G}$. In the case of ultrafilters, this reduces to the Rudin-Keisler ordering, because, between ultrafilters, $\supseteq$ implies $=$.

Remark 15. The conclusion of the proposition says that, between some "large" sets in X and Y, f is a bijection. So for most practical purposes, we can act as if there is a bijection between X and Y sending $\mathcal{U}$ to $\mathcal{V}$. In this case, we say that $\mathcal{U}$ and $\mathcal{V}$ are isomorphic, written $\mathcal{U} \cong \mathcal{V}$.

Another important construction of ultrafilters is as limits of other ultrafilters. Special cases include sums and tensor products.

Definition 16. Let $\{U_i\}_{i\in I}$ be an indexed family of ultrafilters on a set X, and let $\mathcal{V}$ be an ultrafilter on the index set I. The *limit* of the $\mathcal{U}_i$'s with respect to $\mathcal{V}$ is the ultrafilter

$$\mathcal{V}\text{-}\lim_i \mathcal{U}_i = \{A \subseteq X : \{i \in I : A \in \mathcal{U}_i\} \in \mathcal{V}\}$$

on X.

In terms of quantifiers, the definition says that

$$((\mathcal{V}\text{-}\lim_i \mathcal{U}_i)x)\, \varphi(x) \iff (\mathcal{V}i)(\mathcal{U}_i x)\, \varphi(x),$$

which is easily seen to respect propositional connectives, so $\mathcal{V}\text{-}\lim_i \mathcal{U}_i$ is an ultrafilter. Also, in terms of limits of families in compact Hausdorff spaces, limits of ultrafilters yield iterated topological limits:

$$(\mathcal{V}\text{-}\lim_i \mathcal{U}_i)\text{-}\lim_x p_x = \mathcal{V}\text{-}\lim_i (\mathcal{U}_i\text{-}\lim_x p_x)$$

Finally, as the terminology suggests, $\mathcal{V}\text{-}\lim_i \mathcal{U}_i$ is the limit, in the topological sense, of the points $\mathcal{U}_i$ along the ultrafilter $\mathcal{V}$ in the Stone-Čech compactification βX.

An important special case of this notion of limit occurs when the family of ultrafilters $\mathcal{U}_i$ is *strongly discrete*, meaning that they contain sets $A_i \in \mathcal{U}_i$ with $A_i \cap A_j = \varnothing$ for $i \neq j$. The prototypical example of this is when they are ultrafilters on $I \times Y$, with each $\mathcal{U}_i$ containing the corresponding "fiber" $\{i\} \times Y$. In this case, one usually works with the obvious copies on Y of these ultrafilters, and one uses the following terminology and notation.

Definition 17. Let $\{\mathcal{U}_i\}_{i\in I}$ be an indexed family of ultrafilters on Y, and let $\mathcal{V}$ be an ultrafilter on the index set I. The *sum* of the $\mathcal{U}_i$ with respect to $\mathcal{V}$ is the ultrafilter

$$\mathcal{V}\text{-}\sum_i \mathcal{U}_i = \{A \subseteq I \times Y : \{i \in I : \{y \in Y : \langle i, y\rangle \in A\} \in \mathcal{U}_i\} \in \mathcal{V}\}.$$

Equivalently, $\mathcal{V}\text{-}\sum_i \mathcal{U}_i$ can be described as $\mathcal{V}\text{-}\lim_i m_i(\mathcal{U}_i)$, where $m_i : Y \to I \times Y$ is the injection of the i^{th} fiber, $m_i(y) = \langle i, y\rangle$. In terms of quantifiers,

$$((\mathcal{V}\text{-}\sum_i \mathcal{U}_i)\,\langle i,y\rangle)\varphi(i,y) \iff (\mathcal{V}i)(\mathcal{U}_i y)\,\varphi(i,y).$$

Notice that the two projections from $I \times Y$ to I and to Y send $\mathcal{V}\text{-}\sum_i \mathcal{U}_i$ to $\mathcal{V}$ and to $\mathcal{V}\text{-}\lim_i \mathcal{U}_i$, respectively.

In the even more special case where all of the $\mathcal{U}_i$ are the same ultrafilter $\mathcal{U}$, we obtain the tensor product (sometimes called the Fubini product).

Definition 18. Let $\mathcal{V}$ and $\mathcal{U}$ be ultrafilters on I and Y, respectively. Then their *tensor product* $\mathcal{V} \otimes \mathcal{U}$ is the ultrafilter

$$\mathcal{V} \otimes \mathcal{U} = \{A \subseteq I \times Y : \{i \in I : \{y \in Y : \langle i,y\rangle \in A\} \in \mathcal{U}\} \in \mathcal{V}\}.$$

Thus, $\mathcal{V} \otimes \mathcal{U} = \mathcal{V}\text{-}\sum_i \mathcal{U}$. In terms of quantifiers,

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

Now consider the case $X = I = \omega$. We observe that for any two non-principal ultrafilters $\mathcal{U}, \mathcal{V}$ on ω, there are at least two ultrafilters on $\omega \times \omega$ whose projection to the first coordinate is $\mathcal{V}$ and whose projection to the second coordinate is $\mathcal{U}$, namely $\mathcal{V} \otimes \mathcal{U}$ and the reflection of $\mathcal{U} \otimes \mathcal{V}$ across the diagonal . We can tell these are distinct ultrafilters since the former contains the set of $\langle x,y\rangle$ such that $y > x$ and the latter contains the set of $\langle x,y\rangle$ such that $x > y$.

Theorem 19 (Puritz [31])**.** *If $\mathcal{V}$,$\mathcal{U}$ are nonprincipal ultrafilters on ω, then $\mathcal{V} \otimes \mathcal{U}$ is the only ultrafilter $\mathcal{W}$ on ω^2 such that the first projection of $\mathcal{W}$ is $\mathcal{V}$, the second is $\mathcal{U}$, and for any $f : \omega \to \omega$ that is not constant on any set in $\mathcal{U}$, the set of $\langle a,b\rangle$ such that $a < f(b)$ is in $\mathcal{W}$.*

Limits have the following "associativity" property, which is easily proved by just writing out the definitions in full (preferably in terms of quantifiers).

Proposition 20. *Let $\{\mathcal{U}_j : j \in J\}$ be a J-indexed family of ultrafilters on a set X, let $\{\mathcal{V}_i : i \in I\}$ be an I-indexed family of ultrafilters on J, and let $\mathcal{W}$ be an ultrafilter on I. Then*

$$(\mathcal{W}\text{-}\lim_i \mathcal{V}_i)\text{-}\lim_j \mathcal{U}_j = \mathcal{W}\text{-}\lim_i(\mathcal{V}_i\text{-}\lim_j \mathcal{U}_j).$$

A theorem of Mary Ellen Rudin says that, if we consider strongly discrete families of ultrafilters on ω, then this associativity is essentially the only way for two limits to coincide. Here is the precise statement of the result.

Theorem 21 (Rudin [33]). *Assume that* $\{\mathcal{V}_i : i \in \omega\}$ *and* $\{\mathcal{V}'_j : j \in \omega\}$ *are strongly discrete families of ultrafilters on* ω*, that* $\mathcal{U}$ *and* $\mathcal{U}'$ *are ultrafilters on* ω*, and that* $\mathcal{U}\text{-}\lim_i \mathcal{V}_i = \mathcal{U}'\text{-}\lim_j \mathcal{V}'_j$*. Then one of the following three situations occurs.*

- *There are ultrafilters* $\mathcal{W}_i$ *such that* $\mathcal{U}' = \mathcal{U}\text{-}\lim_i \mathcal{W}_i$ *and, for* $\mathcal{U}$*-most* i*,* $\mathcal{V}_i = \mathcal{W}_i\text{-}\lim_j \mathcal{V}'_j$*.*
- *There are ultrafilters* $\mathcal{W}_j$ *such that* $\mathcal{U} = \mathcal{U}'\text{-}\lim_j \mathcal{W}_j$ *and, for* $\mathcal{U}'$*-most* j*,* $\mathcal{V}'_j = \mathcal{W}_j\text{-}\lim_i \mathcal{V}_i$*.*
- $\mathcal{U} \cong \mathcal{U}'$ *via some isomorphism* f *and, for* $\mathcal{U}$*-most* i*,* $\mathcal{V}_i = \mathcal{V}'_{f(i)}$*.*

Note that the first two of the three alternatives here make the assumed equation $\mathcal{U}\text{-}\lim_i \mathcal{V}_i = \mathcal{U}'\text{-}\lim_j \mathcal{V}'_j$ an instance of Proposition 20.

4. Ultraproducts

A frequently useful way to view ultrafilters and study their properties is as "the things you use to form ultraproducts". Consider an X-indexed family of structures, all for the same first-order language, say $\mathfrak{A}_x = (A_x, R^i_x, F^j_x)$ for $x \in X$. Here each R^i_x is the interpretation in $\mathfrak{A}_x$ of the relation symbol R^i, and similarly for function symbols F^j. Given an ultrafilter $\mathcal{U}$ on X, we define the *ultraproduct* $\mathcal{U}\text{-prod}_x\,\mathfrak{A}_x$ as follows. Its underlying set is obtained from the product $\prod_{x\in X} A_x$ by identifying two elements f and g of this product when $(\mathcal{U}x)\,(f(x) = g(x))$; we write $[f]_\mathcal{U}$ or simply $[f]$ for the equivalence class of f. The relation symbols are interpreted by (using a binary relation as a typical example)

$$R^i([f],[g]) \iff (\mathcal{U}x)\, R^i_x(f(x), g(x)).$$

Similarly, the function symbols are interpreted by

$$F^j([f],[g]) = [h] \quad \text{where } h(x) = F^j_x(f(x), g(x)) \text{ for each } x \in X.$$

Łoś's theorem states that, for any formula $\varphi(u, v)$ (say with two free variables for notational simplicity),

$$(\mathcal{U}\text{-prod}_x\,\mathfrak{A}_x) \models \varphi([f],[g]) \iff (\mathcal{U}x)\,(\mathfrak{A}_x \models \varphi(f(x), g(x))).$$

This is essentially built into the definition of the ultraproduct in the case of atomic φ, and the general case is proved by induction on φ.

When $\mathcal{U}$ is principal, say generated by $\{m\}$, the ultraproduct construction just returns (up to isomorphism) the factor indexed by m; the isomorphism $\mathcal{U}\text{-prod}_x\,\mathfrak{A}_x \to \mathfrak{A}_m$ sends each $[f]$ to $f(m)$.

If all the structures $\mathfrak{A}_x$ are the same $\mathfrak{A}$, we write the ultraproduct as $\mathcal{U}\text{-prod}\,\mathfrak{A}$ and call it the *ultrapower* of $\mathfrak{A}$ with respect to $\mathcal{U}$. Łoś's theorem implies that the canonical embedding $\mathfrak{A} \to \mathcal{U}\text{-prod}\,\mathfrak{A}$, sending each $a \in A$ to the equivalence class of the constant function with value

a, is an elementary embedding. We sometimes use it to identify $\mathfrak{A}$ with an elementary substructure of the ultrapower.

Conversely, given any method to produce an elementary extension of an arbitrary structure, we obtain a method of producing an ultrafilter on any set. For any set X, consider the structure $\mathfrak{X}$ consisting of X and all relations and functions on X. If $\mathfrak{X} \preceq \mathfrak{Y}$, then each $y \in \mathfrak{Y}$ determines an ultrafilter on X by $\{A \subseteq X : \mathfrak{Y} \vDash A(y)\}$. This ultrafilter is called the *type* of y in $\mathfrak{Y}$. The ultrapower of $\mathfrak{X}$ with respect to the type of y embeds elementarily into $\mathfrak{Y}$ by the map $[f] \mapsto {}^*f(y)$, where ${}^*f$ is the interpretation in $\mathfrak{Y}$ of the function symbol that denotes in $\mathfrak{X}$ the function $f : X \to X$.

We record, in the following two propositions, how the ultrapower construction interacts with some of the methods in the preceding section for producing new ultrafilters. The proofs are straightforward verifications using the definitions and Łoś's theorem. In the first proposition, we use the notation $\cong\prec$ to mean "elementary embedding", i.e., isomorphism to an elementary submodel.

Proposition 22. *Let $\mathcal{U}$ be an ultrafilter on X, f a function $X \to Y$, and $\{\mathfrak{A}_y : y \in Y\}$ a Y-indexed family of structures, all for the same language. Then there is an elementary embedding*

$$f(\mathcal{U})\text{-}\mathit{prod}_y\, \mathfrak{A}_y \cong\prec \mathcal{U}\text{-}\mathit{prod}_x\, \mathfrak{A}_{f(x)}$$

sending any $[g]_{f(\mathcal{U})}$ to $[g \circ f]_{\mathcal{U}}$.

Proposition 23. *Let $\{\mathcal{U}_i : i \in I\}$ be an indexed family of ultrafilters on a set Y, $\mathcal{V}$ an ultrafilter on the index set I, and $\{\mathfrak{A}_{i,y} : i \in I, y \in Y\}$ and $I \times Y$-indexed family of structures, all for the same language. Then*

$$(\mathcal{V}\text{-}\sum_i \mathcal{U}_i)\text{-}\mathit{prod}_{\langle i,y\rangle}\, \mathfrak{A}_{i,y} \cong \mathcal{V}\text{-}\mathit{prod}_i\, (\mathcal{U}_i\text{-}\mathit{prod}_y\, \mathfrak{A}_{i,y})$$

via an isomorphism sending any $[g]_{\mathcal{V}\text{-}\sum_i \mathcal{U}_i}$ to $[\hat{g}]_{\mathcal{V}}$ where $\hat{g}(i) = [y \mapsto g(\langle i, y\rangle)]_{\mathcal{U}_i}$ for each $i \in I$.

Corollary 24. *Let $\mathcal{V}$ and $\mathcal{U}$ be ultrafilters on I and Y, respectively, and let $\{\mathfrak{A}_{i,y}\}$ be a family of structures, all for the same language, indexed by $I \times Y$. Then*

$$(\mathcal{V} \otimes \mathcal{U})\text{-}\mathit{prod}_{\langle i,y\rangle}\, \mathfrak{A}_{i,y} \cong \mathcal{V}\text{-}\mathit{prod}_i\, (\mathcal{U}\text{-}\mathit{prod}_y\, \mathfrak{A}_{i,y}).$$

Applying the corollary to the case of ultrapowers, where all the structures $\mathfrak{A}_{i,y}$ are the same $\mathfrak{A}$, we find that

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

In this case, the two projection maps from $I \times Y$ to I and to Y induce, by Proposition 22, elementary embeddings of the ultrapowers $\mathcal{V}$-prod $\mathfrak{A}$ and $\mathcal{U}$-prod $\mathfrak{A}$ into $(\mathcal{V} \otimes \mathcal{U})$-prod $\mathfrak{A}$.

For the rest of this section, we specialize to the case of ultrafilters on ω and ultrapowers of the standard model $\mathfrak{N}$ of full arithmetic. Here $\mathfrak{N}$ is the structure with underlying set ω and with all finitary relations and functions on ω. Notice that, because the language has symbols for all relations and functions on ω, every formula is equivalent to an atomic formula, and, between elementary extensions of $\mathfrak{N}$, every embedding is elementary.

If $\mathcal{U}$ is any ultrafilter on ω, the ultrapower $\mathcal{U}$-prod $\mathfrak{N}$ is generated by a single element, the equivalence class [id] of the identity function on ω. Indeed, each element $[f]$ in $\mathcal{U}$-prod $\mathfrak{N}$ is ${}^*f([\mathrm{id}])$.

Conversely, if an elementary extension $\mathfrak{A}$ of $\mathfrak{N}$ is generated by a single element a, then it is isomorphic to the ultrapower $\mathcal{U}$-prod $\mathfrak{N}$, where $\mathcal{U}$ is the type of a in $\mathfrak{A}$ and where the isomorphism sends any $[f] \in \mathcal{U}$-prod $\mathfrak{N}$ to ${}^*f(a) \in \mathfrak{A}$.

Furthermore, any finitely generated extension of $\mathfrak{N}$ is generated by a single element and is therefore isomorphic to an ultrapower of $\mathfrak{N}$. The reason is the availability of pairing functions in the language of $\mathfrak{N}$; these allow any finite number of elements to be coded by a single element. It is sometimes convenient, though, to forgo the coding and work with more than one generator. For example, if a, b are two elements of an extension $\mathfrak{A}$ of $\mathfrak{N}$, then they determine an ultrafilter on ω^2, their *type*, namely $\mathcal{U} = \{X \subseteq \omega^2 : {}^*X(a, b)\}$, and the submodel they generate is isomorphic to $\mathcal{U}$-prod $\mathfrak{N}$. Under the isomorphism, $[f]_{\mathcal{U}}$ corresponds to ${}^*f(a, b)$; in particular, the two projections $\omega^2 \to \omega$ correspond to a and b.

Using these ideas, we can reformulate Theorem 19 as follows.

Corollary 25. *Suppose a and b are elements of some $\mathfrak{C} \succeq \mathfrak{N}$, generating submodels $\mathfrak{A}$ and $\mathfrak{B}$, respectively. Suppose further that all the nonstandard elements of $\mathfrak{B}$ are greater than all elements of $\mathfrak{A}$. If $\mathcal{V}$ and $\mathcal{U}$ are the types of a and b, respectively, then the type of the pair a, b is $\mathcal{V} \otimes \mathcal{U}$.*

Elementary extensions of $\mathfrak{N}$ are structured into constellations: Two elements a and b of such a model are in the same *constellation* if they generate the same submodel. Equivalently, a and b are in the same constellation if and only if $a = {}^*f(b)$ for some one-to-one function f.

A coarser partition of any elementary extension of $\mathfrak{N}$ is given by its skies: Two elements a and b are in the same *sky* if they generate the same initial segment submodel (i.e., an initial segment via the ordering

${}^*\leq$ that corresponds to the standard ordering of the natural numbers in $\mathfrak{N}$). Since the downward closure of any submodel is also a submodel, this is equivalent to $a \leq b \leq {}^*f(a)$ for some f or vice versa. Skies are order-convex; that is, if a sky contains a and b then it also contains every c that is between them in the sense of ${}^*<$. Therefore, the set of skies in any model has a linear ordering induced by ${}^*<$. This ordering of the skies corresponds to the inclusion ordering among the initial segment submodels.

The top sky of $\mathcal{U}$-prod $\mathfrak{N}$ is the set of $[f]_\mathcal{U}$ such that f is finite-to-one. More generally, for any element a in any elementary extension of $\mathfrak{N}$, every element ${}^*f(a)$ in the submodel generated by a is either in the same sky as a or in an earlier sky; it is in the same sky if and only if f is finite-to-one on some set in the type of a.

Given two elementary extensions of $\mathfrak{N}$, we can amalgamate them, i.e., embed both of them elementarily into another such model. In addition, one can specify exactly which parts of the two models are to be identified in the amalgamation. That is, if $\mathfrak{A}$ and $\mathfrak{B}$ are elementary extensions of $\mathfrak{N}$, and if $\theta : \mathfrak{A}' \cong \mathfrak{B}'$ is an isomorphism between submodels $\mathfrak{A}' \preceq \mathfrak{A}$ and $\mathfrak{B}' \preceq \mathfrak{B}$, then there is a model $\mathfrak{C}$ with elementary embeddings $\alpha : \mathfrak{A} \to \mathfrak{C}$ and $\beta : \mathfrak{B} \to \mathfrak{C}$ such that $\alpha(a) = \beta(\theta(a))$ for all $a \in \mathfrak{A}'$ but $\alpha(\mathfrak{A} - \mathfrak{A}')$ is disjoint from $\beta(\mathfrak{B} - \mathfrak{B}')$. In other words, $\mathfrak{A}'$ and $\mathfrak{B}'$ are identified, along the given isomorphism θ, but nothing else is identified.

Furthermore, one can also specify arbitrarily the relative ordering of the skies of $\mathfrak{A}$ above $\mathfrak{A}'$ and the skies of $\mathfrak{B}$ above $\mathfrak{B}'$. (For this and related results about amalgamation of ultrapowers of $\mathfrak{N}$, see [4].) What one cannot do, though, is to map such a "high" sky of $\mathfrak{A}$ and a "high" sky of $\mathfrak{B}$ into the same sky of $\mathfrak{C}$. This last statement is a consequence of the next proposition, applied to elements $a \in \alpha(\mathfrak{A} - \mathfrak{A}')$ and $b \in \beta(\mathfrak{B} - \mathfrak{B}')$. We state the proposition explicitly and give its proof, because the idea behind the proof occurs quite frequently in the theory of ultrafilters on ω.

Proposition 26. *If a, b are in the same sky, then there exist finite-to-one $p, q : \omega \to \omega$ such that ${}^*p(a) = {}^*q(b)$. Equivalently, if f, g are finite-to-one on a set in $\mathcal{U}$, there exist finite-to-one p, q with $p \circ f = q \circ g$ on a set in $\mathcal{U}$.*

Proof. Recall first that any two elements a, b of an elementary extension of $\mathfrak{N}$ lie in a submodel isomorphic to an ultrapower of $\mathfrak{N}$, namely the submodel generated by a code c for the pair $\langle a, b\rangle$. Furthermore, the initial segment submodel generated by the larger of a and b also contains the smaller of the two and therefore the pair code c; that is, the generator c of our submodel is in the same sky as the larger of a and b. In the situation of the proposition, where a and b are in the same sky, then that sky is the top sky of the ultrapower model generated by c. Therefore, a and b are represented by finite-to-one functions in that ultrapower. This shows that the first assertion in the proposition follows from the second, so we prove the second.

Suppose, without loss of generality, that f and g are finite-to-one everywhere (the proof is easily modified otherwise). Partition ω into finite intervals that are so long that, for all x, $f(x)$ and $g(x)$ are always in the same interval or adjacent intervals. That is, define the first interval arbitrarily. Then, if $f(x)$ is in the first interval and $g(x)$ is not, make the second interval long enough to contain $g(x)$, and vice versa. (This is always possible since f and g are finite-to-one.) Continue.

Color the intervals with three colors, say black, green, red, in a cyclically repeating pattern of length 3. Observe that each of the ultrafilters $f(\mathcal{U})$ and $g(\mathcal{U})$ contains the union of all intervals of one of the colors, say black and red, respectively. (If both have the same color, the argument is even easier.) Make a new partition of ω into intervals, coarser than the previous one by making cuts only in green intervals. $f^{-1}(\text{black})$ and $g^{-1}(\text{red})$ are in $\mathcal{U}$, and so their intersection is also in $\mathcal{U}$. If x is in this intersection, then f and g map to adjacent black or red blocks of the first partition, and therefore to the same interval of the second partition. Take $p = q$ to be constant on the intervals of the second partition. $\square$

Note that Proposition 26 implies that, in any ultrapower of $\mathfrak{N}$, the intersection of any two cofinal submodels is cofinal.

5. Special Ultrafilters

In this section, we describe some ultrafilters with especially nice properties. Unless stated otherwise, we assume that we are dealing with non-principal ultrafilters on ω, but the definitions and results could be transferred along a bijection to ultrafilters on any countably infinite set.

Definition 27. An ultrafilter $\mathcal{U}$ is *selective* if whenever ω is partitioned into pieces not in $\mathcal{U}$, there is an $A \in \mathcal{U}$ such that A meets each piece in at most one point. Equivalently, any $f : \omega \to \omega$ becomes either constant or 1-1 when restricted to some set in $\mathcal{U}$.

The second version of the definition amounts to saying that the non-standard part of the ultrapower $\mathcal{U}$-prod $\mathfrak{N}$ is a single constellation. It also shows that selective ultrafilters are exactly the minimal non-principal ultrafilters under the Rudin-Keisler ordering.

Theorem 28 (Kunen, published in [14]). *If $\mathcal{U}$ is a selective ultrafilter on ω, then, for any partition of $[\omega]^2$ into two pieces, there is a homogeneous set in $\mathcal{U}$, i.e., a set $H \in \mathcal{U}$ such that $[H]^2$ is included in one of the pieces. Furthermore, the same is true for any finite number of pieces and for partitions of $[\omega]^k$ for any finite k.*

Proof of the first assertion. Identifying $[\omega]^2$ with the above-diagonal part $\{\langle x, y\rangle : x < y\}$ of ω^2, we find that the statement to be proved is equivalent to the statement that the filter $\mathcal{H}$ on ω^2 generated by $[\omega]^2$ and the sets $A \times A$ for $A \in \mathcal{U}$ is an ultrafilter. Clearly $\mathcal{U} \otimes \mathcal{U}$ is an ultrafilter extending $\mathcal{H}$; we must show (in the light of Proposition 5) that it is the only ultrafilter extending $\mathcal{H}$.

Consider, therefore, any ultrafilter $\mathcal{W}$ extending $\mathcal{H}$, and consider the ultrapower $\mathcal{W}$-prod $\mathfrak{N}$. It is generated by $\langle a, b\rangle$, where a and b are the equivalence classes of the two projections $\omega^2 \to \omega$. Because $\mathcal{W} \supseteq \mathcal{H}$, we know that a and b both have type $\mathcal{U}$ and that $a < b$. Consider the submodels $\mathfrak{A}$ and $\mathfrak{B}$ generated by a and b, respectively, in $\mathcal{W}$-prod $\mathfrak{N}$. They are copies of $\mathcal{U}$-prod $\mathfrak{N}$, and so they each have, because $\mathcal{U}$ is selective, a single constellation of non-standard elements.

If these two constellations, one from $\mathfrak{A}$ and one from $\mathfrak{B}$, are in different skies of $\mathcal{W}$-prod $\mathfrak{N}$, then all the nonstandard elements of $\mathfrak{B}$ are above all elements of $\mathfrak{A}$. Then by Puritz's result, Corollary 25, $\mathcal{W} = \mathcal{U} \otimes \mathcal{U}$, as desired.

There remains the case that the non-standard constellations of $\mathfrak{A}$ and $\mathfrak{B}$, which contain a and b, lie in the same sky of $\mathcal{W}$-prod $\mathfrak{N}$. In this case, Proposition 26 provides an element c that is simultaneously of the forms ${}^*p(a)$ and ${}^*q(b)$ for some finite-to-one $p, q : \omega \to \omega$. This c is therefore in both $\mathfrak{A}$ and $\mathfrak{B}$, and, as it is non-standard, it generates each of these submodels. That is, $\mathfrak{A}=\mathfrak{B}$. In particular, each of a and b is in the submodel generated by the other, so $a = {}^*f(b)$ for some $f : \omega \to \omega$. Since both a and b have type $\mathcal{U}$, we obtain $f(\mathcal{U}) = \mathcal{U}$, and therefore, by Corollary 14, f is the identity on some set in $\mathcal{U}$. That implies ${}^*f(b) = b$, i.e., $a = b$, contrary to the fact that $a < b$. $\square$

The extension of this result to partitions with more than two pieces is easy. The theorem provides homogeneous sets in $\mathcal{U}$ for all the two-piece coarsenings of a given partition, and we need only intersect these homogeneous sets.

The extension to partitions of $[\omega]^k$ for $k > 2$ requires more work, but it can be done similarly to the proof for $k = 2$. It is based on studying the possible amalgamations of k copies of $\mathcal{U}$-prod $\mathfrak{N}$ (in place of the two copies $\mathfrak{A}$ and $\mathfrak{B}$ in the preceding proof). See [5] and the references there for more information about this and related proofs.

The partition property in the first assertion of Theorem 28 easily implies selectivity. Given a function f on ω, partition $[\omega]^2$ by putting $\{x, y\}$ into one piece if and only if $f(a) = f(b)$. On a homogeneous set, f is one-to-one or constant.

Theorem 28 adds to Ramsey's theorem the information that the homogeneous sets can always be found in any prescribed selective ultrafilter. Because selectivity is equivalent to this strengthening of Ramsey's theorem, selective ultrafilters are often called *Ramsey* ultrafilters.

Ramsey's theorem cannot be extended to arbitrary partitions of the set $[\omega]^\omega$ of infinite subsets of ω. The axiom of choice easily yields a partition of $[\omega]^\omega$ into two pieces such that, when two infinite sets differ by a single element, they are in different pieces. Such a partition cannot have a homogeneous set. Nevertheless, an infinitary partition relation holds when the partition is sufficiently well-behaved. Recall that we topologize $[\omega]^\omega$ as a subspace of a product 2^ω of discrete spaces; recall also that a set is called analytic if it is the image of a Borel set under a continuous mapping. Then we have the following infinitary partition theorem for selective ultrafilters.

Theorem 29 (Mathias [26])**.** *If $\mathcal{U}$ is selective and if $[\omega]^\omega$ is partitioned into an analytic and a co-analytic piece, then there is a homogeneous set $H \in \mathcal{U}$.*

In specific models of set theory, the requirement of analyticity in this theorem can be relaxed, but, as indicated above, it cannot be removed altogether (unless one discards the axiom of choice). Mathias has shown that this partition property holds for all partitions that are ordinal definable from reals in the model obtained by collapsing to ω all cardinals below some Mahlo cardinal.

The existence of selective ultrafilters cannot be proved in ZFC; there are none in the random real model [23]. It follows, however, from CH, or from Martin's Axiom, or from $\mathfrak{c} = \mathbf{cov}(\mathcal{B})$ (the minimum number of meager, or first-category sets needed to cover the real line, also the minimum number of closed sets with empty interiors needed to cover the real line). In fact, it is shown in [17] that, under the last of these assumptions (which is the weakest of the three), any filter on ω that contains the cofinite sets and is generated by fewer than $\mathfrak{c}$ sets can be extended to a selective ultrafilter.

Forcing with (the separative quotient of) the partial order $([\omega]^\omega, \subseteq)$ produces a selective ultrafilter without adding new reals. It can be used to prove, from Mathias's theorem, the earlier result of Silver [36], not involving ultrafilters in its statement, that is like Mathias's theorem except that H is merely required to be infinite, not in any prescribed ultrafilter.

An important weakening of selectivity is the following, which was studied even before selectivity [34].

Definition 30. A nonprincipal ultrafilter $\mathcal{U}$ is a *P-point* if every $f : \omega \to \omega$ becomes either constant or finite-to-one when restricted to some set in $\mathcal{U}$.

An equivalent characterization is that, for any countably many sets $A_n \in \mathcal{U}$, there is some $B \in \mathcal{U}$ that is almost included in each A_n, i.e., $B - A_n$ is finite for all n. Such a B is called a *pseudo-intersection* of the A_n's. This notion of P-point is the specialization, to the space $\beta\omega - \omega$ of non-principal ultrafilters on ω, of the general topological notion of P-point, namely a point x such that, for any countably many neighborhoods of x, there is a single neighborhood included in them all.

Kunen's Theorem 28 admits a generalization for P-points $\mathcal{U}$. For the case of exponent 2, it says that, for each partition of $[\omega]^2$ into finitely many pieces, there exist a set $H \in \mathcal{U}$ and a function $f : \omega \to \omega$ such that one piece of the partition contains all pairs $\{a < b\} \in [H]^2$ for which $f(a) < b$. For higher exponents, the statement is similar; we get homogeneity for those subsets of H whose elements are "far apart" as measured by f.

$\mathcal{U}$ is a P-point if and only if $\mathcal{U}$-prod $\mathfrak{N}$ has only one sky of non-standard elements.

The existence of P-points, in fact of P-points that are not selective, follows from CH, or just Martin's Axiom. ZFC alone does not prove the existence of P-points.

Definition 31. A non-principal ultrafilter $\mathcal{U}$ on ω is a *Q-point* if every function that is finite-to-one on a set in $\mathcal{U}$ is one-to-one on a set in $\mathcal{U}$.

Thus, an ultrafilter is selective if and only if it is both a P-point and a Q-point.

$\mathcal{U}$ is a Q-point if and only if the top sky of $\mathcal{U}$-prod $\mathfrak{N}$ is a single constellation.

The existence of Q-points follows from CH, or just Martin's Axiom, but it is not provable in ZFC. There exist models of ZFC without P-points [35, 40] and models of ZFC without Q-points [29], but it is an open problem whether there exists a model without either.

6. Cardinal Characteristics of the Continuum

In this section, we introduce some of the many cardinal characteristics of the continuum, especially those that have interesting connections with the theory of ultrafilters. For a general survey of cardinal characteristics, see [9, 19, 39].

In this context, "the continuum" could mean $\mathbb{R}$, Cantor space 2^ω, Baire space ω^ω, $[\omega]^\omega$, etc. These spaces are essentially the same, in that for any pair, after removal of at most a countable set from each space, there exists a homeomorphism between the modified spaces. For those that carry natural measures, the homeomorphism can be taken to preserve measure as well, provided the total measures of the pair are equal. Thus, the ambiguity of the terminology "continuum" will cause no real problems.

The idea of a cardinal characteristic is that for some combinatorial property, $\aleph_0$ and $\mathfrak{c}$ behave differently. We can look at the least cardinal which behaves like $\mathfrak{c}$, and this is called the cardinal characteristic associated to that property. This is, of course, uninteresting if CH holds, for then all such cardinal characteristics are equal to $\mathfrak{c}$.

Definition 32. In this definition, we work in ω^ω.

(1) f *dominates* g if, for all but finitely many $n \in \omega$, $f(n) \geq g(n)$.
(2) The *dominating number* $\mathfrak{d}$ is the minimum cardinality of a dominating family – a subset of ω^ω such that every g is dominated by an f in the family.
(3) The *unbounding number* $\mathfrak{b}$ is the minimum cardinality of an unbounded family – a subset of ω^ω not dominated by a single function.

It is easily provable in ZFC that $\aleph_1 \leq \mathfrak{b} \leq \mathfrak{d} \leq \mathfrak{c}$. Also, $\mathfrak{b}$ is a regular cardinal and $\mathfrak{b} \leq \mathrm{cf}(\mathfrak{d})$.

Definition 33. In this definition, we work in $[\omega]^\omega$.

(1) A set $S \subseteq \omega$ *splits* an infinite $A \subseteq \omega$ if both $A \cap S$ and $A - S$ are infinite.
(2) The *splitting number* $\mathfrak{s}$ is the minimum size of a splitting family – a family $\mathcal{S} \subseteq [\omega]^\omega$ such that for all $A \in [\omega]^\omega$, there is $S \in \mathcal{S}$ that splits A.

(3) The *unsplitting (or refining or reaping) number* $\mathfrak{r}$ is the minimum cardinality of an unsplit family – a family $\mathcal{R} \subseteq [\omega]^\omega$ such that no single set splits all the sets in $\mathcal{R}$.

Definition 34. In this definition, we work with subsets of the real line.

(1) The *covering number* for Baire category $\mathbf{cov}(\mathcal{B})$ is the minimum cardinality of a family of meager sets whose union covers the real line. Define $\mathbf{cov}(\mathcal{L})$ in the same way, except with "(Lebesgue) measure 0" in place of meager.[2]
(2) The *cofinality* of category $\mathbf{cof}(\mathcal{B})$ is the minimum cardinality of a basis for the ideal $\mathcal{B}$ of meager sets, i.e., a family of meager sets such that every meager set is included in one from the family. Again, the definition is similar for $\mathbf{cof}(\mathcal{L})$, replacing meager with measure 0.

Remark 35. All of these cardinal charaacteristics can be expressed as the minimum cardinality of a family of reals such that, for a certain relation, every real is related to a real in the family. (In the case of meager or measure zero sets, which are not themselves reals, we can restrict attention to F_σ or G_δ sets, respectively, and these can be coded by reals.) Note that $\mathfrak{b}$ and $\mathfrak{d}$ can be obtained from each other by taking the negation of the converse of the specified relation. The same is true for $\mathfrak{s}$ and $\mathfrak{r}$. In this sense, these pairs of cardinals are dual to each other. The covering and cofinality numbers also have duals, called the uniformity and additivity numbers, respectively.

Theorem 36 (Ketonen [22])**.** *If $\mathfrak{d} = \mathfrak{c}$, then there exist P-points.*

Theorem 37 (Canjar [17])**.** *If $\mathbf{cov}(\mathcal{B}) = \mathfrak{d}$, then there exist Q-points.*

Cichoń's diagram is a diagram showing the ZFC-provable inequalities between ten cardinal characteristics. For any ideal $\mathcal{J}$ of sets (in particular for the ideal $\mathcal{B}$ of meager sets and the ideal $\mathcal{L}$ of measure-zero sets) the *additivity* $\mathbf{add}(\mathcal{J})$ is the dual of $\mathbf{cof}(\mathcal{J})$, i.e., the minimum cardinality of a family of sets in $\mathcal{J}$ whose union is not in $\mathcal{J}$. The *uniformity* $\mathbf{non}(\mathcal{J})$ is the dual of $\mathbf{cov}(\mathcal{J})$, i.e., the minimum cardinality of a set not in $\mathcal{J}$. An arrow in the diagram from a cardinal $\mathfrak{k}$ to a cardinal $\mathfrak{l}$ means that ZFC proves $\mathfrak{k} \leq \mathfrak{l}$. Every inequality between two of these cardinals which is provable in ZFC is represented by an arrow or a sequence of arrows.

[2]We use $\mathcal{B}$ and $\mathcal{L}$, in honor of Baire and Lebesgue, to denote the ideals of meager and measure-zero sets, respectively. Other notations for $\mathcal{B}$ include $\mathcal{K}$ (for "Kategorie") and $\mathcal{M}$ (for "meager"); other notations for $\mathcal{L}$ include $\mathcal{M}$ (for "measure") and $\mathcal{N}$ (for "null").

$$\begin{tikzcd}
\mathbf{cov}(\mathcal{L}) \arrow[r] & \mathbf{non}(\mathcal{B}) \arrow[r] & \mathbf{cof}(\mathcal{B}) \arrow[r] & \mathbf{cof}(\mathcal{L}) \\
 & \mathfrak{b} \arrow[u] \arrow[r] & \mathfrak{d} \arrow[u] & \\
\mathbf{add}(\mathcal{L}) \arrow[uu] \arrow[r] & \mathbf{add}(\mathcal{B}) \arrow[u] \arrow[r] & \mathbf{cov}(\mathcal{B}) \arrow[u] \arrow[r] & \mathbf{non}(\mathcal{L}) \arrow[uu]
\end{tikzcd}$$

The only further restrictions on the values of these cardinals are $\mathbf{add}(\mathcal{B}) = \min\{\mathfrak{b}, \mathbf{cov}(\mathcal{B})\}$ and $\mathbf{cof}(\mathcal{B}) = \max\{\mathfrak{d}, \mathbf{non}(\mathcal{B})\}$. In fact, for any assignment of $\aleph_1$ and $\aleph_2$ to these ten cardinals which respects the diagram and these two additional restrictions, there is a model of ZFC where each cardinal has the assigned value; see [3, Chapter 7].

Additionally, $\mathfrak{r}$ and $\mathfrak{s}$ interact nicely with the cardinals in Cichoń's diagram. In ZFC, $\mathfrak{r} \geq \mathbf{cov}(\mathcal{L}), \mathfrak{b}, \mathbf{cov}(\mathcal{B})$, and $\mathfrak{s} \leq \mathbf{non}(\mathcal{L}), \mathfrak{d}, \mathbf{non}(\mathcal{B})$.

7. Cardinal Characteristics and Ultrafilters

We consider ultrafilters on ω, and as usual assume $\mathcal{U}$ is nonprincipal. The cardinality of $\mathcal{U}$ is $\mathfrak{c}$. Instead of asking about the cardinality of $\mathcal{U}$, we can ask how many sets does it take to generate $\mathcal{U}$ (i.e., by closing under finite intersections and supersets). This number is called the *character* of $\mathcal{U}$ and denoted by $\chi(\mathcal{U})$. Equivalently, it is the minimum cardinality of a base for $\mathcal{U}$. It is easy to see that $\chi(\mathcal{U})$ is always between $\aleph_1$ and $\mathfrak{c}$. In fact, it is at least $\mathfrak{r}$, since an ultrafilter base has to be an unsplit family.

Definition 38. $\mathfrak{u}$ is the minimum $\chi(\mathcal{U})$ over all non-principal ultrafilters $\mathcal{U}$ on ω.

By the remarks above, we always have $\mathfrak{r} \leq \mathfrak{u}$. Goldstern and Shelah [20] have shown that $\mathfrak{r} < \mathfrak{u}$ is consistent with ZFC. However, Aubrey [1] has shown that $\mathfrak{r} \geq \min\{\mathfrak{u}, \mathfrak{d}\}$, so if $\mathfrak{r} < \mathfrak{u}$ then $\mathfrak{d}$ has to be smaller than $\mathfrak{u}$.

While $\mathfrak{u}$ concerns small characters of ultrafilters, one might also ask about big characters. That question is easily answered in ZFC.

Proposition 39. *There is an ultrafilter $\mathcal{U}$ on ω such that $\chi(\mathcal{U}) = \mathfrak{c}$.*

Proof. Use the theorem of Hausdorff from the proof of Proposition 6 above, obtaining an independent family $\mathcal{F}$ of size $\mathfrak{c}$. Let $\mathcal{U}$ be an ultrafilter containing all the sets in $\mathcal{F}$ and the complements of all intersections of infinitely many sets from $\mathcal{F}$. It is easy to check that these sets have the finite intersection property, so such an ultrafilter exists. If $\mathcal{U}$ had a base of cardinality $< \mathfrak{c}$, then, since every $A \in \mathcal{F}$ contains a set from the base, infinitely many sets from $\mathcal{F}$ would have to contain the same B from the base. But then B is included in the intersection of infinitely many sets from $\mathcal{F}$, and so cannot be in $\mathcal{U}$, a contradiction. □

Returning to $\mathfrak{u}$, we have a result due to Solomon [37] that $\mathfrak{b} \leq \mathfrak{u}$: Given a base for an ultrafilter, consider, for each set in the base, the function that moves n to the least $m \geq n$ in this basis set. The functions obtained in this way form an unbounded family. Actually, the argument shows $\mathfrak{b} \leq \mathfrak{r}$, because all that is used about ultrafilter bases is that they are unsplit families. The proof shows further that, if $\mathcal{F}$ is a filter containing all cofinite sets and generated by fewer than $\mathfrak{b}$ sets, then there exists a partition of ω into finite intervals such that each set in the filter meets all but finitely many intervals. By applying the function f that maps all elements of the n^{th} interval to n, we obtain that $f(\mathcal{F})$ is the filter of cofinite sets.

Filters $\mathcal{F}$ with the property that $f(\mathcal{F})$ is the cofinite filter for some finite-to-one function f are called *feeble* filters. A theorem of Talagrand [38] says that these are exactly the filters that are meager as subsets of $\mathcal{P}(\omega) \cong 2^\omega$.

Definition 40. A π-base of an ultrafilter $\mathcal{U}$ is a family $\mathcal{B} \subseteq [\omega]^\omega$ such that for all $X \in \mathcal{U}$, there exists $B \in \mathcal{B}$ such that $B \subseteq X$. The π-*character* of $\mathcal{U}$, written $\pi\chi(\mathcal{U})$, is the smallest cardinality of a π-base of $\mathcal{U}$.

Note that the definition of "π-base" differs from that of "base" only in that a π-base for $\mathcal{U}$ need not be a subset of $\mathcal{U}$.

Theorem 41 (Balcar and Simon [2])**.** *The minimum value of $\pi\chi(\mathcal{U})$ over all non-principal ultrafilters on ω is $\mathfrak{r}$.*

Proof sketch. The minimum in question is at least $\mathfrak{r}$ because a π-base for an ultrafilter is necessarily an unsplit family.

Toward proving the reverse inequality, let us consider under what circumstances a family $\mathcal{B} \subseteq [\omega]^\omega$ can be a π-base for an ultrafilter. We need that, if a set X has no subset in $\mathcal{B}$, then X cannot be in the ultrafilter, and so $\omega - X$ must be in the ultrafilter. That is, the family $\mathcal{C}$ of all complements of the sets not in $\mathcal{B}$ must have the finite intersection property, so that it can be extended to an ultrafilter. Untangling the negation and complement, we can state this requirement more simply: Whenever ω is partitioned into finitely many pieces, at least one piece must have a subset in $\mathcal{B}$.

An unsplit family $\mathcal{R}$ of size $\mathfrak{r}$ has this property for partitions into two pieces. We handle more pieces by iterating. Obtain $\mathcal{R}'$ by puting a copy of $\mathcal{R}$ into each set from $\mathcal{R}$; then $\mathcal{R}'$ has the desired property for partitions into 4 or fewer pieces. Repeat this idea ω times. □

Another cardinal number naturally associated to an ultrafilter $\mathcal{U}$ is the cofinality of the ultrapower, cf($\mathcal{U}$-prod $\mathfrak{N}$).

Theorem 42 (Canjar [15], Roitman [32])**.** *It is consistent with ZFC that $\mathfrak{b} \ll \mathfrak{d}$ and for every regular cardinal $\kappa \in [\mathfrak{b}, \mathfrak{d}]$, there is $\mathcal{U}$ with cf($\mathcal{U}$-prod $\mathfrak{N}$) $= \kappa$.*

The model used for the proof is obtained by adding sufficiently many Cohen reals to a model of CH.

Theorem 43 (Canjar [16])**.** *There is a non-principal ultrafilter $\mathcal{U}$ on ω with cf($\mathcal{U}$-prod $\mathfrak{N}$) $=$ cf($\mathfrak{d}$).*

8. Groupwise Density

The groupwise density number $\mathfrak{g}$ is a modification of a more familiar characteristic, the distributivity number $\mathfrak{h}$. We define them together to emphasize the similarity.

Definition 44.
(1) A family $\mathcal{H} \subseteq [\omega]^\omega$ is *dense* if it is closed under subsets and every set in $[\omega]^\omega$ has a subset in $\mathcal{H}$.
(2) The *distributivity number* $\mathfrak{h}$ is the minimum cardinality of a set of dense families with empty intersection.
(3) A family $\mathcal{G} \subseteq [\omega]^\omega$ is *groupwise dense* if it is closed under subsets and, whenever ω is partitioned into finite intervals, some union of these intervals is in $\mathcal{G}$.
(4) The *groupwise density number* $\mathfrak{g}$ is the minimum cardinality of a set of groupwise dense families with empty intersection.

Remark 45. The distributivity number owes its name to the fact that it is the smallest cardinal κ for which the forcing $([\omega]^\omega, \supseteq)$ is not κ-distributive, i.e., the smallest κ for which this forcing adjoins a new function from κ into the ground model.

Every groupwise dense family is dense: Given an infinite set $X \subseteq \omega$, apply the definition of groupwise density to a partition into intervals such that each interval contains an element of X. Any infinite union of intervals from this partition has a subset that is an infinite subset of X.

It follows immediately that $\mathfrak{h} \leq \mathfrak{g}$. It is also easy to verify that $\mathfrak{h} \leq \mathfrak{b}, \mathfrak{s}$.

Theorem 46 (Blass and Mildenberger [12])**.** *For all non-principal ultrafilters $\mathcal{U}$ on ω, we have $\mathfrak{g} \leq$ cf($\mathcal{U}$-prod $\mathfrak{N}$).*

Theorem 42 can be interpreted as saying that $\mathfrak{b}$ and $\mathfrak{d}$ are, respectively, the only lower and upper bounds on the cofinalities of ultrapowers $\mathcal{U}$-prod $\mathfrak{N}$. The new lower bound $\mathfrak{g}$ is, in some models, strictly larger than $\mathfrak{b}$, but this does not conflict with Theorem 42 because, in the Cohen model used there, $\mathfrak{g} = \mathfrak{b} = \aleph_1$.

The three cardinals $\mathfrak{b}$, $\mathfrak{g}$, and $\mathfrak{s}$ often behave similarly; they are next to each other in diagrams of cardinal characteristics like that in [9, Section 11]. Since both $\mathfrak{b}$ and $\mathfrak{g}$ are lower bounds for cofinalities of ultraproducts, it is reasonable to ask whether $\mathfrak{s}$ is also a lower bound. (It is $\aleph_1$ in the Cohen model, so this would not contradict Theorem 42.) This is probably false[3] although it was shown in [12] that at most one cardinal below $\mathfrak{s}$ can be equal to some $cf(\mathcal{U}\text{-prod}\mathfrak{N})$. Also, at most one cardinal above $\mathfrak{r}$ can equal some $cf(\mathcal{U}\text{-prod}\mathfrak{N})$. So if $\mathfrak{s} > \mathfrak{r}$, then at most two cardinals occur as cofinalities of such ultrapowers.

The definition of $\mathfrak{g}$ was inspired by three statements which are rather difficult to satisfy; most of the "usual" models of set theory don't satisfy them. The first of these three statements is *near coherence of filters* (NCF), which says that, for any two filters $\mathcal{F}$ and $\mathcal{G}$ on ω (containing all the cofinite sets), there exist finite-to-one functions f, g such that their images are coherent – $f(\mathcal{F}) \cup g(\mathcal{G})$ has the finite intersection property (i.e., generates a filter). Equivalently, we need only consider this for ultrafilters (coherence is harder to achieve for bigger filters), and once we've done this, we can replace "coherent" with "equal." Also, we can assume that $f = g$, and that f is nondecreasing. By a theorem of Shelah [13], NCF is consistent with ZFC.

NCF is equivalent to each of the following two statements.

- For every $\mathcal{U}$, there is a finite-to-one f such that $\chi(f(\mathcal{U})) < \mathfrak{d}$. [6]
- $\mathfrak{u} < \min\{cf(\mathcal{U}\text{-prod}\mathfrak{N}) : \mathcal{U}$ a non-principal ultrafilter on $\omega\}$. [27]

Remark 47. Any two ultrafilters with bases of size less than $\mathfrak{d}$ are nearly coherent. This is proved by an interval argument similar to the ones we've seen before.

The second statement that led up to the definition of $\mathfrak{g}$ is *filter dichotomy* (FD), which says that, if $\mathcal{F}$ is any filter on ω containing the cofinite sets, then there is a finite-to-one f such that $f(\mathcal{F})$ is either just

[3] A proof that it is false, by Blass and Mildenberger [12], depends on an earlier result by Blass and Shelah [13, Section 6], which has an error. Shelah and Mildenberger have, independently, given corrected proofs of that result, but this work is not published yet. The proof by Blass and Mildenberger is otherwise correct.

the filter of cofinite sets or an ultrafilter. It is easy to see that FD implies NCF, by taking $\mathcal{F}$ in FD to be the intersection of two ultrafilters. The consistency of FD was proved by Laflamme [11], by modifying one lemma in Shelah's proof of the consistency of NCF.

The third statement is the cardinal characteristic inequality $\mathfrak{u} < \mathfrak{g}$ introduced in [11]. It (and $\mathfrak{g}$ itself) were found by looking at the consistency proof of NCF minus the lemma that Laflamme modified. The inequality $\mathfrak{u} < \mathfrak{g}$ encapsulates the result of that argument, and one can plug in various substitutes for the omitted lemma to get various consequences of $\mathfrak{u} < \mathfrak{g}$.

We have the implications

$$\mathfrak{u} < \mathfrak{g} \implies \text{FD} \implies \text{NCF}.$$

Mildenberger and Shelah [28] have proved that the second of these implications is not reversible; whether the first is reversible remains an open problem.

An equivalent formulation of $\mathfrak{u} < \mathfrak{g}$ is *semifilter trichotomy*, which says that, for every family $\mathcal{F} \subseteq [\omega]^\omega$, if $\mathcal{F}$ is closed upward (under $\subseteq$) and closed under finite changes, then there exists a finite-to-one f such that $f(\mathcal{F})$ is either the filter of cofinite sets, or an ultrafilter, or $[\omega]^\omega$. One direction of the equivalence between semifilter trichotomy and $\mathfrak{u} < \mathfrak{g}$ was proved in [24] and the other direction in [8].

For more information about $\mathfrak{u} < \mathfrak{g}$ and its consequences, see [7] and the references there.

Remark 48. In keeping with the pedagogical character of this paper, the following bibliography occasionally omits original sources in favor of general surveys or books. For the original sources not given here, see the bibliographies of these secondary sources. Good general references for ultrafilters and cardinal characteristics include [18] and [9, 19, 39], respectively.

References

[1] Jason Aubrey, "Combinatorics for the dominating and unsplitting numbers," *J. Symbolic Logic* 69 (2004) 482–498.

[2] Bohuslav Balcar and Petr Simon, "On minimal π-character of points in extremally disconnected compact spaces," *Topology Appl.* 41 (1991) 133–145.

[3] Tomek Bartoszyński and Haim Judah, *Set Theory: On the Structure of the Real Line,* A K Peters (1995).

[4] Andreas Blass, "Amalgamation of nonstandard models of arithmetic," *J. Symbolic Logic* 42 (1977) 372–386.

[5] Andreas Blass, "A model-theoretic view of some special ultrafilters," in *Logic Colloquium '77* (L. Pacholski, A. Macintyre, and J. Paris, eds.) North-Holland (1978) 79–90.

[6] Andreas Blass, "Near coherence of filters, I: Cofinal equivalence of models of arithmetic," *Notre Dame J. Formal Logic* 27 (1986) 579–591.

[7] Andreas Blass, "Applications of superperfect forcing and its relatives," in *Set Theory and Its Applications* (J. Steprāns and S. Watson, eds.) Springer-Verlag, Lecture Notes in Mathematics 1401 (1989) 18–40.

[8] Andreas Blass, "Groupwise density and related cardinals," *Arch. Math. Logic* 30 (1990) 1–11.

[9] Andreas Blass, "Combinatorial cardinal characteristics of the continuum," in *Handbook of Set Theory,* vol. 1 (M. Foreman and A. Kanamori, eds.) Springer-Verlag (2010) 395–489.

[10] Andreas Blass, "Ultrafilters and set theory," in *Ultrafilters Across Mathematics* (V. Bergelson, A. Blass, M. Di Nasso, and R. Jin, eds.) *Contemp. Math. 530* (2010) 49–71.

[11] Andreas Blass and Claude Laflamme, "Consistency results about filters and the number of inequivalent growth types," *J. Symbolic Logic* 54 (1989) 50–56.

[12] Andreas Blass and Heike Mildenberger, "On the cofinality of ultrapowers," *J. Symbolic Logic* 64 (1999) 727–736.

[13] Andreas Blass and Saharon Shelah, "There may be simple $P_{\aleph_1}$ and $P_{\aleph_2}$ points and the Rudin-Keisler ordering may be downward directed," *Ann. Pure Appl. Logic* 33 (1987) 213–243.

[14] David Booth, "Ultrafilters on a countable set," *Ann. Math. Logic* 2 (1970) 1–24.

[15] R. Michael Canjar, *Model-Theoretic Properties of Countable Ultraproducts Without the Continuum Hypothesis,* Ph.D. thesis, University of Michigan (1982).

[16] R. Michael Canjar, "Cofinalities of countable ultraproducts: the existence theorem," *Notre Dame J. Formal Logic* 30 (1989) 539–542.

[17] R. Michael Canjar, "On the generic existence of special ultrafilters," *Proc. Amer. Math. Soc.* 110 (1990) 233–241.

[18] W. Wistar Comfort and Stelios Negrepontis, *Theory of Ultrafilters,* Springer-Verlag, Grundlehren der Mathematischen Wissenschaften 211 (1974).

[19] Eric van Douwen, "The integers and topology," in *Handbook of Set Theoretic Topology* (K. Kunen and J. Vaughan, eds.) North-Holland (1984) 111–167.

[20] Martin Goldstern and Saharon Shelah, "Ramsey ultrafilters and the reaping number, Con($\mathfrak{r} < \mathfrak{u}$)," *Ann. Pure Appl. Logic* 49 (1990) 121–142.

[21] Hewitt Kenyon, "Problem 5077," *Amer. Math. Monthly* 70 (1963) 216.

[22] Jussi Ketonen, "On the existence of P-points," *Fund. Math.* 92 (1976) 91–99.

[23] Kenneth Kunen, "Some points in βN," *Proc. Cambridge Phil. Soc.* 78 (1980) 385–398.

[24] Claude Laflamme, "Equivalence of families of functions on the natural numbers," *Trans. Amer. Math. Soc.* 330 (1992) 307–319.

[25] F. William Lawvere, "Functorial concepts of complexity for finite automata" *Theory and Applications of Categories* 13 (2004) 164–168.

[26] Adrian R. D. Mathias, "Happy families," *Ann. Math. Logic* 12 (1977) 59–111.

[27] Heike Mildenberger, "Groupwise dense families," *Arch. Math. Logic* 40 (2001) 93–112.

[28] Heike Mildenberger and Saharon Shelah, "The near coherence of filters principle does not imply the filter dichotomy principle," *Trans. Amer. Math. Soc.* 361 (2009) 2305–2317.

[29] Arnold Miller, "There are no Q-points in Laver's model for the Borel conjecture," *Proc. Amer. Math. Soc.* 78 (1980) 103–106.

[30] Bedřich Pospíšil, "Remark on bicompact spaces," *Ann. of Math. (2),* 38 (1937) 845–846.

[31] Christian Puritz, "Skies, constellations and monads," in *Contributions to Non-Standard Analysis* (W. A. J. Luxemburg and A. Robinson, eds.) North-Holland, Studies in Logic and Foundations of Math. 69 (1972) 215–243.

[32] Judith Roitman, "Non-isomorphic H-fields from non-isomorphic ultrapowers," *Math. Z.* 181 (1982) 93–96.

[33] Mary Ellen Rudin, "Types of ultrafilters," *1966 Topology Seminar (Wisconsin, 1965)*, Princeton Univ. Press, Ann. of Math. Studies 60, 147–151.

[34] Walter Rudin, "Homogeneity problems in the theory of Čech compactifications," *Duke Math. J.* 23 (1956) 409–419.

[35] Saharon Shelah, *Proper Forcing,* Springer-Verlag, Lecture Notes in Mathematics 940 (1982).

[36] Jack Silver, "Every analytic set is Ramsey," *J. Symbolic Logic* 35 (1970) 60–64.

[37] R. C. Solomon, "Families of sets and functions," *Czechoslovak Math. J.* 27 (1977) 556–559.

[38] Michel Talagrand, "Compacts de fonctions mesurables et filtres non mesurables," *Studia Math.* 67 (1980) 13–43.

[39] Jerry Vaughan, "Small uncountable cardinals and topology," in *Open Problems in Topology* (J. van Mill and G. Reed, eds.) North-Holland (1990) 195–218.

[40] Edward Wimmers, "The Shelah P-point independence theorem," *Israel J. Math.* 43 (1982) 28–48.

Mathematics Department, University of Michigan, Ann Arbor, MI 48109–1043, U.S.A.

E-mail address: ablass@umich.edu