Notes for the Tenth Appalachian Set Theory Workshop
From the lectures given by Andreas Blass
On the topics of ultrafilters and cardinal characteristics of the continuum
Delivered on Jan. 30, 2010 at the Miami University in Oxford, Ohio
Compiled by Nicholas Rupprecht

1. Ultrafilters

Definition 1. An ultrafilter on the set X is a $\mathcal{U} \subseteq \mathcal{P}(X)$ such that

(1) $\varnothing \notin \mathcal{U}$,
(2) $X \in \mathcal{U}$,
(3) $A \subseteq B$ and $A \in \mathcal{U}$ implies $B \in \mathcal{U}$,
(4) $A, B \in \mathcal{U}$ implies $A \cap B \in \mathcal{U}$,
(5) For all $A \subseteq X$, $A \in \mathcal{U}$ or $X - A \in \mathcal{U}$,
(6) $A \cup B \in \mathcal{U}$ implies $A \in \mathcal{U}$ or $B \in \mathcal{U}$.

Remark 2. Some of these are redundant. The first four alone define a filter on X. For (4) and (6), the converses are also true by (3). (5) may also be read as $A \in \mathcal{U} \iff X - A \notin \mathcal{U}$. We can think of this as: membership in $\mathcal{U}$ respects boolean combinations (propositional connectives). In this way, we can view $\mathcal{U}$ as a map from 2^X to 2 such that for any operation on 2^k (and the operation it induces on $(2^X)^k$), $\mathcal{U}$ commutes with the operations. That is, we think of $\mathcal{U}$ as mapping f to the i such that $f^{-1}(i) \in \mathcal{U}$ – we define an ultrafilter in this way iff $\mathcal{U}$ commutes with all k-ary operations as stated above. Further, it is enough that $\mathcal{U}$ commutes with binary and unary operations (as commuting with k-ary operations for $k > 2$ can be inferred from commuting with these). An interesting result of Lawvere says that if we instead consider maps from 3^X to 3, any map which commutes with just the unary operations gives an ultrafilter $\mathcal{U}$ (where $f^{-1}(i)$ is in $\mathcal{U}$ iff f is mapped to i).

Another approach is to view ultrafilters as quantifiers:

$$(\mathcal{U}x)\phi(x) \iff \{x \in X : \phi(x)\} \in \mathcal{U}$$

In a sense, $(\mathcal{U}x)\phi(x)$ means for "almost all" x, $\phi(x)$ holds. The quantifier respects propositional connectives.

A third approach is that an ultrafilter represents a uniform way to choose limits of sequences. $\mathcal{U}$ amounts to an operation assigning to each X-indexed family $\{p_x\}_{x \in X}$ of points in a compact Hausdorff space C a single point $\mathcal{U}\text{-}\lim_x p_x \in C$ in a way that commutes with continuous functions: if $f : C \to C'$ is continuous,

$$\mathcal{U}\text{-}\lim_x f(p_x) \in C' = f(\mathcal{U}\text{-}\lim_x p_x \in C)$$

Intuitively, $\mathcal{U}$ selects a point such that every neighborhood of it contains "almost all" p_x.

A fourth approach is to let $\mathcal{U}$ be a point in βX, the Stone-Čech compactification of X.

Given $\mathcal{U}$ an ultrafilter on X and structures $\mathfrak{A}_x = (A_x, R^i_x, F^j_x)$, the ultrapower $\mathcal{U}\text{-prod}_x \mathfrak{A}_x$ is the structure with universe

$$A = \prod_{x \in X} A_x$$

modulo $f \sim g$ if $(\mathcal{U}x) f(x) = g(x)$ and relations and functions defined by $R^i([f])$ iff $(\mathcal{U}x) R^i_x(f(x))$, and so on. Łoś's theorem states that for all formulas ϕ, ϕ holds in $\mathcal{U}\text{-prod}_x \mathfrak{A}_x$ iff it is true in $\mathcal{U}$ many $\mathfrak{A}_x$ (though in the definition we only guarantee this for atomic ϕ). In particular, if $\mathfrak{A}_k = \mathfrak{A}$ for all k, we get an embedding of $\mathfrak{A}$ into $\mathcal{U}\text{-prod}_x \mathfrak{A}_x$ (taking $a \in \mathfrak{A}$ to $[f]_{\mathcal{U}}$ where $f(x) = a$ for all x), and the theorem says it is an elementary embedding.

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 sets and functions on X. If $\mathfrak{X} \preceq \mathfrak{Y}$ each $y \in \mathfrak{Y}$ determines an ultrafilter on X by $\{A \subseteq X : \mathfrak{Y} \models A(y)\}$ (the type of y in $\mathfrak{Y}$).

A principal (trivial) ultrafilter is given by $A \in \mathcal{U} \iff x \in A$ for a particular $x \in X$. Note it is principal in the sense that it is generated by a single element. $\mathcal{U}$ is principal iff it contains a singleton iff it contains a finite set. So $\mathcal{U}$ is non-principal iff it contains all cofinite sets.

For principal $\mathcal{U}$, the interpretation in the other senses is either simply evaluating at x, taking x in βX, or the ultraproduct structure is $\mathfrak{A}_x$. In each case, we are doing something trivial (hence, we may call principal ultrafilters trivial).

By Zorn's lemma, every filter on X is contained in some ultrafilter on X. Hence, there exist non-principal ultrafilters on X if X is infinite (start with the filter of cofinite sets). Also, the ultrafilters are exactly the maximal filters. Another way to think of this proof is to well-order the subsets of X, and proceed by transfinite recursion along this ordering. At the step labelled by a particular set, decide (if it hasn't already been decided) whether to put the set or its complement into $\mathcal{U}$. At the end, we have $\mathcal{U}$ is an ultrafilter. Since we make $2^{|X|}$ choices, it makes sense that:

Proposition 3. *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 many subsets of X such that given any disjoint finite subfamilies $\mathcal{A}, \mathcal{B}$, the intersection of sets in $\mathcal{A}$ and complements of sets in $\mathcal{B}$ is nonempty.

Hausdorff's example begins by saying it is enough to do this for

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

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

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

So for any choice of "positive" Y's and "negative" Y's, we want (P, Q) such that the intersection of P with the Y's is in or not in Q (respectively). Choose P such that these intersections are all different.

Now, given a family $\mathcal{F}$ (of subsets of X), for each $\mathcal{G} \subseteq \mathcal{F}$, there is an ultrafilter containing the sets from $\mathcal{G}$ and the complements of the sets from $\mathcal{F} - \mathcal{G}$ (Hausdorff's result says there is at least a filter for which this is true). □

We show a use of the quantifier view of ultrafilters by proving (without much work) a weak version of Ramsey's theorem. 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 there is an $H \subseteq \omega$ such that $|H| = \alpha$ and $[H]^k \subseteq$ one piece.

Theorem 4. *For any n, $\omega \to (n)^k_c$.*

Proof. View k-element subsets as increasing k-tuples. Denote the pieces of the partition by $C_1, \dots, C_c$. Then

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

Fix a nonprincipal ultrafilter $\mathcal{U}$. We can replace each $\forall$ with $\mathcal{U}$, and the statement is still true. Now push the disjunction over i to the outside, and fix an i such that the inner statement is true. Rewrite this statement by renaming variables and introducing dummy variables as follows. For any subset of $\{1, \dots, n\}$ of size k ($\{r_1, \dots, r_k\}$), let x_l be replaced with y_{r_l}, and let the remaining y_j be dummy variables. We can push this conjunction (over all choices for $\{r_1, \dots, r_k\}$) inside the ultrafilter quantifiers. Replace $\mathcal{U}y_1$ with $\exists y_1$, $\mathcal{U}y_2$ with $\exists y_2 > y_1$, etc. A witness to this statement gives the desired H. □

We can tackle a harder theorem (similar to one of Nash-Williams):

Theorem 5. *If you partition $[\omega]^\omega$ into an open piece and a closed piece (under the topology induced from 2^ω), then there is an infinite $H \subseteq \omega$ such that $[H]^\omega$ is in one piece.*

Proof. Again, instead of thinking of infinite subsets, we think of infinite increasing sequences. So $[\omega]^\omega$ is identified as the set of paths through $\omega^{\uparrow<\omega}$ (i.e., this tree only contains finite increasing sequences). The open piece can be specified by a set of nodes in this tree. Choose a nonprincipal ultrafilter $\mathcal{U}$ on ω. Mark nodes by ordinals, as follows. Mark 0 on nodes for the open set. For $\alpha > 0$, mark α on a node s iff $(\mathcal{U}n)s^\frown\langle n\rangle$ has been marked with a $\beta < \alpha$.

Case 1: The empty sequence gets marked by α for some α.

Pick 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, $\mathcal{U}$ many h_n will work, and there are only finitely many F. Any infinite subset of the h_i is in the open piece (along this infinite subset, the marks keep decreasing, so must hit 0 eventually).

Case 2: The empty sequence is not marked.

Note that if a node is unmarked, $\mathcal{U}$ many of its 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 the h_i, all initial segments are unmarked (in particular, not marked with 0), so the set is in the closed piece. □

Remark 6. From this, we can prove, for any k, that $\omega \to (\omega)^k_2$. Given a partitioning of $[\omega]^k$ into C_0 and C_1, partition $[\omega]^\omega$ into D_0 and D_1, where $H \in D_i$ iff the set of the first k elements of H is in C_i. The pieces D_0 and D_1 are clopen, giving an infinite H such that for some i $[H]^\omega \subseteq D_i$, and hence $[H]^k \subseteq C_i$.

If $\mathcal{U}$ is an ultrafilter on X, and $f : X \to Y$, there is an ultrafilter on Y denoted $f(\mathcal{U})$:

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

Another view is that $(f(\mathcal{U})y)\phi(y)$ is equivalent to $(\mathcal{U}x)\phi(f(x))$, or that $f(\mathcal{U})$ is the evaluation at $\mathcal{U}$ of the unique continuous extension of f to the Stone-Čech compactifications. Also, the $f(\mathcal{U})$ ultraproduct of $\mathfrak{A}_y$ is canonically elementarily embedded into the $\mathcal{U}$ ultraproduct of $\mathfrak{A}_{f(x)}$.

This induces the Rudin-Keisler ordering: $\mathcal{V} \leq_{RK} \mathcal{U}$ if $\mathcal{V} = f(\mathcal{U})$ for some f. This is reflexive and transitive, but not antisymmetric (so not a partial order). However:

Proposition 7. *If $\mathcal{V} \leq_{RK} \mathcal{U} \leq_{RK} \mathcal{V}$, there is an f such that $f(\mathcal{U}) = \mathcal{V}$ and f is 1-1 on a set in $\mathcal{U}$.*

Proof. We use the following general lemma.

Lemma 8. *If $f : X \to X$, there is 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 $i = 1, 2, 3$, $f(A_i) \cap A_i = \varnothing$.*

Let f and g witness $\mathcal{V} \leq_{RK} \mathcal{U}$ and $\mathcal{U} \leq_{RK} \mathcal{V}$, respectively. Then $f \circ g$ is a map from Y to itself. Let A_i for $i = 0, 1, 2, 3$ be as in the lemma. Note that $\mathcal{V}$ must contain some A_i, and also $f \circ g(A_i)$. We must have $i = 0$. Now note that $g(A_0) \in \mathcal{U}$ and f restricted to this set is 1-1. □

Remark 9. That is, on 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. In this case, we say $\mathcal{U} \cong \mathcal{V}$. We note that in this case, the elementary embedding of ultraproducts mentioned above is an isomorphism.

Now assume ultrafilters are on ω. From a $\mathcal{U}$ we get an ultrapower of $\mathfrak{N}$, which is ω with all relations and functions. If $\mathcal{U} \leq_{RK} \mathcal{V}$, then the $\mathcal{U}$ ultrapower of $\mathfrak{N}$ is canonically elementarily embedded in the $\mathcal{V}$ ultrapower of $\mathcal{N}$. If f witnesses $\mathcal{U} \leq_{RK} \mathcal{V}$, then $[g]_{\mathcal{U}}$ is mapped to $[g \circ f]_{\mathcal{V}}$. Any such ultrapower is generated by a single element ($[\mathrm{id}]_{\mathcal{U}}$) – that is, the only submodel containing $[\mathrm{id}]_{\mathcal{U}}$ is the entire ultrapower since $[f]_{\mathcal{U}} = {}^*f([\mathrm{id}]_{\mathcal{U}})$. We note that any embedding between two of these elementary extensions of $\mathfrak{N}$ is in fact an elementary embedding.

Elementary extensions of $\mathfrak{N}$ are structured into constellations: a, b are in the same constellation iff they generate the same submodel. Equivalently, $a = {}^*f(b)$ for some 1-1 f. We can also structure into skies: a, b are in the same sky iff they generate the same initial segment submodel (i.e., an initial segment via the ordering ${}^* \leq$). 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. The skies are naturally linearly ordered. The top sky of $\mathcal{U}$-prod $\mathfrak{N}$ is the set of $[f]_{\mathcal{U}}$ such that f is finite-to-one.

Given two such ultrapowers, we can amalgamate them, ordering the non-standard skies above the identified parts in any way. However, we cannot identify skies without identifying their elements:

Proposition 10. *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. We prove the second statement. First suppose f and g are finite-to-one everywhere (proof easily modified otherwise). Partition ω into long finite intervals such that for all x, $f(x)$ and $g(x)$ are always in the same interval or adjacent intervals. That is, define the first interval in any way, if $f(x)$ is in the first interval and $g(x)$ is not, make the second interval long enough to contain it, and vice versa (this is always possible since f and g are finite-to-one). Continue.

Color the intervals with three colors (in a repeating pattern of length 3 – black, green, red). Observe that $f(\mathcal{U})$ and $g(\mathcal{U})$ each contains the union of all intervals of one of the colors – without loss of generality, black and red, respectively. Make a new partition coarser than the previous one by making cuts only in green intervals. f^{-1}(black) and g^{-1}(red) are in $\mathcal{U}$, hence, their intersection is also in $\mathcal{U}$. If x is in this intersection, f and g map to adjacent blocks of the first partition, hence in the same interval of the second partition. Take $p = q$ to be constant on the intervals of the second partition. □

Remark 11. Proofs of this type are common. Note that the proposition implies the intersection of two cofinal submodels is cofinal.

Suppose $\{\mathcal{U}_i\}_{i\in I}$ is an indexed family of nonprincipal ultrafilters on some set X, and $\mathcal{V}$ is a nonprincipal ultrafilter on I. Then $\mathcal{V}\text{-}\lim_i \mathcal{U}_i$ is the set of $A \subseteq X$ such that $(\mathcal{V}i)A \in \mathcal{U}_i$, which is an ultrafilter. We have:

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

We can also define $\mathcal{V}\text{-}\Sigma_i\mathcal{U}_i$ as the set of $A \subseteq I \times X$ such that

$$(\mathcal{V}i)(\mathcal{U}_i x)(i,x) \in A.$$

The projection down to I is $\mathcal{V}$ and the projection across to X is $\mathcal{V}\text{-}\lim_i \mathcal{U}_i$. If all the $\mathcal{U}_i = \mathcal{U}$, then denote $\mathcal{V}\text{-}\Sigma_i\mathcal{U}_i$ by $\mathcal{V}\otimes\mathcal{U}$. Note that

$$((\mathcal{V}\otimes\mathcal{U})(i,x))\phi(i,x) \iff (\mathcal{V}i)(\mathcal{U}x)\phi(i,x)$$

Now consider the case $X = I = \omega$. We observe that for any two 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 across the diagonal of $\mathcal{U}\otimes\mathcal{V}$. We can tell these are distinct ultrafilters since the former contains the set of (x,y) such that $y > x$ and the latter contains the set of (x,y) such that $x > y$.

Proposition 12. *(Puritz) 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 $f:\omega\to\omega$ non-constant on any set in $\mathcal{U}$, the set of (a,b) such that $a < f(b)$ is in $\mathcal{W}$.*

In terms of ultrapowers, the $\mathcal{W}$ ultrapower of $\mathfrak{N}$ contains copies of the $\mathcal{V}$ and $\mathcal{U}$ ultrapowers such that all of the $\mathcal{V}$ ultrapower is below the nonstandard part of the $\mathcal{U}$ ultrapower.

We have another fact:

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

Furthermore by a theorem of M. E. Rudin, if we consider ultrafilters on ω, if $\mathcal{U}\text{-}\lim_i \mathcal{V}_i = \mathcal{U}'\text{-}\lim_j \mathcal{V}'_j$ and if the sequences $\{\mathcal{V}_i\}$ and $\{\mathcal{V}'_j\}$ are strongly discrete (i.e., can be covered by a sequence of pairwise disjoint neighborhoods in $\beta\omega$), then either $\mathcal{U}'$ is a $\mathcal{U}$ limit of some $\mathcal{W}_i$ and for $\mathcal{U}$ many i $\mathcal{V}_i$ is a $\mathcal{W}_i$ limit of the $\mathcal{V}'_j$, or vice versa, or $\mathcal{U} \cong \mathcal{U}_i$ via some f and for $\mathcal{U}$ many i, $\mathcal{V}_i = \mathcal{V}'_{f(i)}$. Note the first two cases mean the equality can be expressed as an instance of the iterated limit formula above.

Definition 13. An ultrafilter $\mathcal{U}$ is selective if whenever ω is partitioned into pieces not in $\mathcal{U}$, there is $A \in \mathcal{U}$ 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 $\mathcal{U}$.

Remark 14. Selective ultrafilters are minimal under $\leq_{RK}$.

Theorem 15. *(Kunen) If $\mathcal{U}$ is selective, then for any partition of $[\omega]^2$ into 2 pieces, there is a homogeneous set in $\mathcal{U}$. Further, this is true for any number of pieces and any $[\omega]^k$.*

Remark 16. Equivalently (to the first statement), the filter (on $[\omega]^2$) generated by $A \times A$ for $A \in \mathcal{U}$ and $[\omega]^2$ is an ultrafilter. This ultrafilter must be $\mathcal{U} \otimes \mathcal{U}$. For the ultrapower situation, an ultrapower of $\mathfrak{N}$ via a selective ultrafilter is a minimal nonstandard one (i.e. any nonstandard element generates the whole thing). The theorem can thus be thought of as a consequence of the fact that the amalgamation of two such ultrapowers must have one ultrapower's nonstandard part completely in front of the other's.

Theorem 17. *(Mathias) If $\mathcal{U}$ is selective and you partition $[\omega]^\omega$ into an analytic and a co-analytic piece, there is $H \in \mathcal{U}$ homogeneous.*

The existence of a selestive ultrafilter on ω follows from CH, or from Martin's Axiom (MA), or from $\mathfrak{c} = \mathbf{cov}(\mathcal{B})$ (the minimum number of meager sets, or first category sets, needed to cover the real line, or the minimum number of closed sets without interior needed to cover the real line). In fact, under the latter assumption, for any filter containing the cofinite sets and generated by fewer than $\mathfrak{c}$ sets, that filter can be extended to a selective ultrafilter. The existence of selective ultrafilters does not follow from ZFC alone (adding enough random reals produces a model without selective ultrafilters). On the other hand, forcing with the separative quotient of $([\omega]^\omega, \subseteq)$ produces a model with a selective ultrafilter. In fact, this forcing can be used to prove (from Mathias's theorem) the theorem which removes $\mathcal{U}$ and says H is merely infinite.

Consider the following weakening of selectivity:

Definition 18. 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 $\mathcal{U}$.

CH and MA imply the existence of P-points which are not selective. This definition has an interesting topological equivalent. In $\beta\omega - \omega$, every countably many neighborhoods of $\mathcal{U}$ include a neighborhood of $\mathcal{U}$. Kunen's theorem (for partitions of $[\omega]^2$) has an analogous version for P-points – instead of H being homogeneous, there is an f such that for all $a < b \in H$ with $f(a) < b$, the pair (a, b) gets the same color.

A Q-point is an ultrafilter $\mathcal{U}$ such that any function finite-to-one on a set in $\mathcal{U}$ is 1-1 on a set in $\mathcal{U}$. There exist models of ZFC without P-points and models of ZFC without Q-points, but it is an open problem if there exists a model without either.

2. Cardinal Characteristics of the Continuum

Here, the continuum could mean $\mathbb{R}$, Cantor space ${}^{\omega}2$, Baire space ${}^{\omega}\omega$, $[\omega]^{\omega}$, 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 (which can be assumed to be measure-preserving whenever the spaces have natural measures and both spaces have measure equal to the same value with respect to their individual measures). Therefore, we may refer to all of them as "the continuum."

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 $\mathfrak{c}$ (assuming CH fails – otherwise this is uninteresting).

Definition 19. We work in ${}^{\omega}\omega$.

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

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

Definition 20. (1) $\mathfrak{s}$ is the minimum size of $\mathcal{S} \subseteq [\omega]^\omega$ such that for all $A \in [\omega]^\omega$, there is $S \in \mathcal{S}$ such that $|A \cap S| = |A - S| = \aleph_0$ (we say that such an S splits A).

(2) $\mathfrak{r}$ is the minimum cardinality of $\mathcal{R} \subseteq [\omega]^\omega$ such that for all $A \in \mathcal{P}(\omega)$ there exists $R \in \mathcal{R}$ with $R \subseteq^* A$ or $R \subseteq^* \omega - A$ (that is, the minimum cardinality of a family of infinite sets not split by any single set).

(3) $\mathbf{cov}(\mathcal{B})$ 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.

(4) $\mathbf{cof}(\mathcal{B})$ is the minimum cardinality of a basis for the ideal $\mathcal{B}$ of meager sets (again, definition is similar for $\mathcal{L}$ instead of $\mathcal{B}$, replacing meager with measure 0).

Remark 21. All of these can be expressed as the minimum cardinality of a family of reals such that for some relation, every real is related to a real in the family (for meager or measure zero, we look at F_σ or G_δ sets 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 duals of the other cardinals are also well-known, but not needed for this lecture).

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

Theorem 23. *(Canjar) If $\mathbf{cov}(\mathcal{B}) = \mathfrak{d}$, there exist Q-points.*

Cichoń's diagram is a diagram which relates how 10 cardinals relate to each other (in ZFC). In the diagram below, $\mathbf{add}(\mathcal{J})$ is the dual of $\mathbf{cof}(\mathcal{J})$ (the minimum cardinality of a family of sets in $\mathcal{J}$ whose union is not in $\mathcal{J}$), $\mathbf{non}(\mathcal{J})$ is the dual of $\mathbf{cov}(\mathcal{J})$ (the minimum cardinality of a set not in $\mathcal{J}$), and an arrow from a cardinal $\mathfrak{k}$ to a cardinal $\mathfrak{l}$ means that ZFC proves $\mathfrak{k} \leq \mathfrak{l}$. Any inequality between two cardinals which is provable in ZFC is represented by an arrow or a sequence of arrows.

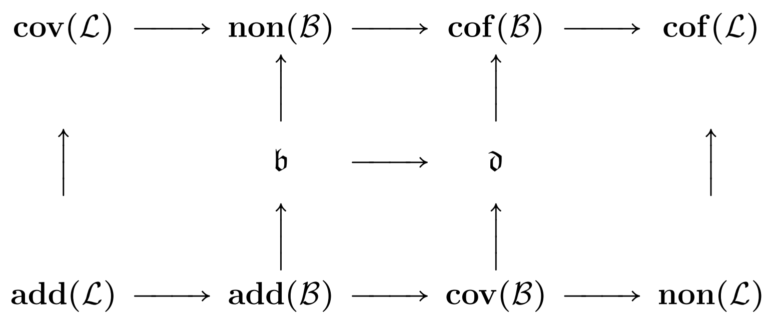

The only further restrictions on the values of these cardinals are $\mathbf{add}(\mathcal{B}) = \min\{\mathfrak{b}, \mathbf{cov}(\mathcal{B})\}$ and $\mathbf{cof}(\mathcal{B}) = \max\{\mathfrak{d}, \mathbf{non}(\mathcal{B})\}$. In fact, for any assignment of $\aleph_1$ and $\aleph_2$ to cardinals which respects the diagram and these additional 2 restrictions, there is a model of ZFC where each cardinality equals its assignment.

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 $\mathfrak{s} \leq \mathbf{non}(\mathcal{L}), \mathfrak{d}, \mathbf{non}(\mathcal{B})$.

3. Connections between ultrafilters and cardinal characteristics

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). Denote it by $\chi(\mathcal{U})$. Equivalently, it is the minimum cardinality of a base for $\mathcal{U}$. Trivially, this is between $\aleph_1$ and $\mathfrak{c}$. In fact, this is at least $\mathfrak{r}$, since the base has to be an unsplit family (the family $\mathcal{R}$ from the definition).

Definition 24. $\mathfrak{u}$ is the smallest value of $\chi(\mathcal{U})$ possible.

So by the above, $\mathfrak{r} \leq \mathfrak{u}$. It is consistent with ZFC (by Goldstern and Shelah) that $\mathfrak{r} < \mathfrak{u}$. However, ZFC proves $\mathfrak{r} \geq \min\{\mathfrak{u}, \mathfrak{d}\}$ (by Aubrey, from his Ph.D. thesis), so if $\mathfrak{r} < \mathfrak{u}$, then $\mathfrak{d}$ is small.

Note $\mathfrak{u}$ talks about small characters. What about big characters? The question is simply answered.

Proposition 25. *There is an ultrafilter $\mathcal{U}$ such that $\chi(\mathcal{U}) = \mathfrak{c}$.*

Proof. Use the theorem of Hausdorff from before, obtaining an independent family $\mathcal{F}$ of size $\mathfrak{c}$. Let $\mathcal{U}$ be an ultrafilter containing all sets in $\mathcal{F}$, and the complement of any infinite intersection of sets in $\mathcal{F}$ (as an exercise, check these sets have the finite intersection property). If $\mathcal{U}$ had a small base, every $A \in \mathcal{F}$ contains a set from the base, so infinitely many sets from $\mathcal{F}$ contain the same B from the base. But then B is in an infinite intersection of sets from $\mathcal{F}$, a contradiction. □

Returning to $\mathfrak{u}$, we have (due to Solomon) $\mathfrak{b} \leq \mathfrak{u}$: given a base, for each basis set, consider the function that moves n to the least $m \geq n$ in the basis set. The functions obtained in this way form an unbounded family. Actually, the argument shows $\mathfrak{b} \leq \mathfrak{r}$. Further, if $\mathcal{F}$ is a filter containing all cofinite sets generated by fewer than $\mathfrak{b}$ elements, 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 a function that maps the nth interval to n, we obtain the filter of cofinite sets (an $\mathcal{F}$ with this property is called feeble; this property is equivalent to being meager).

Definition 26. A π-base of an ultrafilter $\mathcal{U}$ is a $\mathcal{B} \subseteq [\omega]^\omega$ such that for all $X \in \mathcal{U}$, there exists $B \in \mathcal{B}$ such that $B \subseteq X$. $\pi\chi(\mathcal{U})$ is the smallest cardinality of a π-base of $\mathcal{U}$.

Theorem 27. *(Balcar) The minimum cardinality of $\pi\chi(\mathcal{U})$ (over all $\mathcal{U}$) is $\mathfrak{r}$.*

Proof. (Sketch) Suppose you have a family $\mathcal{B}$ of infinite subsets of ω. $\mathcal{B}$ can only be a π-base if a set that doesn't contain any $B \in \mathcal{B}$ is not going to be in the ultrafilter (i.e., its complement is). We need such sets to have the finite intersection property. It turns out that the necessary thing is for any partition of ω into finitely many pieces, some piece contains some $B \in \mathcal{B}$. A family of size $\mathfrak{r}$ gives us this for 2 pieces. We get more pieces by iterating: for each member of the family $\mathcal{R}$, break it into pieces which copy $\mathcal{R}$. Repeat ω many times. □

Definition 28. Let $\mathcal{U}$ be an ultrafilter. We consider $cf(\mathcal{U}\text{-prod}\mathfrak{N})$, the cofinality of the linearly ordered ultrapower of $\mathfrak{N}$. That is, the minimum size of a family $\mathcal{G} \subseteq \omega^\omega$ such that for all $f \in \omega^\omega$, there is $g \in \mathcal{G}$ such that $(\mathcal{U}n) f(n) < g(n)$.

For all $\mathcal{U}$, $\mathfrak{b} \leq cf(\mathcal{U}\text{-prod}\mathfrak{N}) \leq \mathfrak{d}$. In some sense, this is the best we can say.

Theorem 29. *(Canjar, Roitman) It is consistent with ZFC that $\mathfrak{b} \ll \mathfrak{d}$ and for all regular $\kappa \in [\mathfrak{b}, \mathfrak{d}]$, there is $\mathcal{U}$ with $cf(\mathcal{U}\text{-prod}\mathfrak{N}) = \kappa$.*

Remark 30. This is done by adding sufficiently many Cohen reals.

Theorem 31. *(Canjar) There exists $\mathcal{U}$ such that $cf(\mathcal{U}\text{-prod}\mathfrak{N}) = cf(\mathfrak{d})$.*

Definition 32. $\mathfrak{g}$ is the minimum cardinality of a set of groupwise dense families with empty intersection. A groupwise dense family is a $\mathfrak{G} \subseteq [\omega]^\omega$ such that if $X \in \mathfrak{G}$ and $Y \subseteq^* X$, then $Y \in \mathfrak{G}$, and for each partition of ω into finite intervals, some union of these intervals is in $\mathfrak{G}$.

Remark 33. The definition of $\mathfrak{g}$ is a modification of the definition of $\mathfrak{h}$, an older cardinal. $\mathfrak{h}$ is defined in the same way, except delete the word "groupwise", and in the second clause of the second definition, instead say for all infinite X, there exists $Y \subseteq X$ in $\mathfrak{G}$. Easily, any groupwise dense family is dense, so $\mathfrak{h} \leq \mathfrak{g}$. Also, $\mathfrak{h}$ is easily $\leq \mathfrak{b}, \mathfrak{s}$.

Theorem 34. *(Blass and Mildenberger) For all $\mathcal{U}$, $\mathfrak{g} \leq cf(\mathcal{U}\text{-prod}\mathfrak{N})$.*

Remark 35. In the model from the theorem by Canjar and Roitman, $\mathfrak{g} = \aleph_1$, so this gives no further restriction.

One could reasonably ask if $\mathfrak{s} \leq cf(\mathcal{U}\text{-prod}\mathfrak{N})$ also always holds. This is probably false[1] although it is known 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 $\mathfrak{s} > \mathfrak{r}$, all ultrapowers have the same cofinality).

The definition of $\mathfrak{g}$ was inspired by three statements which tend to be false (unless the model is constructed specifically for the statements to be true). One is near coherence of filters. NCF is the statement: for any two filters $\mathcal{F}$ and $\mathcal{G}$ on ω (including the cofinite sets), there exist finite-to-1 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 $f = g$, and f is nondecreasing. By a theorem of Shelah, NCF is consistent with ZFC.

NCF is equivalent to each of:

- For every $\mathcal{U}$, there is a finite-to-1 f such that $\chi(f(\mathcal{U})) < \mathfrak{d}$.
- $\mathfrak{u} < \min\{cf(\mathcal{U}\text{-prod}\mathfrak{N})\}$.

Remark 36. Any two ultrafilters with bases of size less than $\mathfrak{d}$ are nearly coherent. This is done by an interval argument similar to the ones we've seen before.

Another statement which plays a role in the study of $\mathfrak{g}$ is filter dichotomy. FD is the statement: if $\mathcal{F}$ is any filter containing the cofinite sets, there is a finite-to-1 f such that $f(\mathcal{F})$ is either just the filter of cofinite sets or an ultrafilter. It is easy to see that FD implies NCF, and this is also consistent with ZFC – the proof is a modification of a lemma from Shelah's proof, done by Laflamme. If we remove this lemma, (i.e., look at the real result of Shelah's proof) we get $\mathfrak{u} < \mathfrak{g}$ is consistent with ZFC (and implies FD).

Recently, Mildenberger and Shelah proved that NCF does not imply FD, but it is open whether or not FD implies $\mathfrak{u} < \mathfrak{g}$. It is known that $\mathfrak{u} < \mathfrak{g}$ iff 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-1 f such that $f(\mathcal{F})$ is either the filter of cofinite sets, an ultrafilter, or $[\omega]^\omega$.

[1] A proof by Blass and Mildenberger depends on an earlier result by Blass and Shelah, which has an error that may or may not be fixed yet. The proof by Blass and Mildenberger is otherwise correct.