Understanding Galois-Tukey Connections

Tim de Jonge

Universiteit Utrecht

Supervised by:
J. van Oosten

7.5 EC

Department of Mathematics
Utrecht University

17/06/2016

1 Introduction

Usually, our set theory works with ZFC, and in the case of dealing with infinities, we add the Continuum Hypothesis to ensure that if an infinity is bigger than $\aleph_0$, it is at least $\mathfrak{c}$. However, this is far from a necessity - the best known result about infinities is Cantor's diagonal proof that $\mathfrak{c} > \aleph_0$, but between these two infinities lies a relatively undiscovered wasteland, usually hidden by the Continuum Hypothesis (CH). In this thesis, we will answer the question what this wasteland contains, and how we can work with it.

Generally, this is done by statements that have a cardinal number as parameter, where clearly the result is true if the parameter is countable, and false if the parameter is greater than $\mathfrak{c}$. Some examples are as follows:

- κ nowhere dense sets cannot cover the real line.
- If κ sets have Lebesgue measure zero, then so does their union.
- Given κ sequences of real numbers, there is a single sequence that dominates each of the given one.

If κ is countable, these statements are either relatively well-known theorems, or relatively easy to prove. On the other hand, if $\kappa \geq \mathfrak{c}$, we see that these statements are trivially false. So, the question is then, what is the first cardinality for which these statements become true? Here we see that assuming CH, this question is trivial - as there is nothing between $\aleph_0$ and $\mathfrak{c}$, the answer to this question has to be $\mathfrak{c}$. Fortunately, CH is independent from ZFC, and as such, we can take CH to be false. Without CH, the question whether there are other cardinal numbers that satisfy these problems becomes rather more difficult. Specifically so because this question in a vacuum is undecidable in ZFC - if we take $\aleph_2 = \mathfrak{c}$, it could be that $\kappa = \aleph_1$ works for all these questions, or none of them.

The first cardinality for which a certain statement becomes true is also called the *cardinal characteristic* of the statement, and is one of the primary subjects of this thesis. Generally, these cardinal characteristics are used as diagnostic tools to see the effects of a model of ZFC on small infinities. Once a model is formed, we can see to which cardinality these characteristics resolve, and as a result, we can gain some intuition about the model itself.

Now, we cannot prove anything about a single cardinal characteristic in ZFC, but still, we can prove some connections between them. Specifically, under certain circumstances, we can prove inequalities, for example: if κ satisfies the second problem given above, it also satisfies the first.

For this purpose, we will be using Galois-Tukey Connections, which is a relatively simple way to deal with these problems, in a similar way that regularly one would use injective or bijective functions to prove inequalities between infinities. We will then work on establishing some inequalities, and finally, see what this says about a model satisfying Martin's Axiom.

This work has as primary and sole source Blass [1] and is in essence a restructuring, rephrasing, and further explanation of some of the ideas present in his work.

2 Meeting the Characteristics

This section is based on Blass' introduction and chapter 2. [1]

Before we get to work on our first characteristics, we first need some tools. We will initially be dealing with functions ${}^{\omega}\omega$, and specifically, we need a way to extend the notion of inequality to functions. Before we do this though, we need some explanation of what exactly this notation ${}^{\omega}\omega$ means. The set ${}^{\omega}\omega$ is given as the set of functions from the natural numbers onto the natural numbers - ω is the ordinal number associated with the natural numbers.

The most interesting results are created if we work "modulo finite". As such, we change the $\forall$ quantifier to the $\forall^{\infty}$ quantifier: $\forall^{\infty}x$ then means "For all but finitely many x", or equivalently, "for x sufficiently large". To maintain the duality between $\forall$ and $\exists$, we have to define $\exists^{\infty}x$ as "there exist infinitely many x", or "there exist infinitely large x". More commonly used than the ∞ as a superscript is to put an asterisk next to the quantifier or operation. This would mean that $\forall^{\infty}x$ and $\forall^{*}x$ mean the same thing.

Using this notation, the idea of an inequality on ${}^{\omega}\omega$ follows fairly naturally: for $f, g \in {}^{\omega}\omega$, we define $f \leq^{*} g$ if for all but finitely many x we have $f(x) \leq g(x)$.

Note that this would fall just short of being a partial order - we have reflexivity and transitivity, but if $f \leq^{*} g$ and $g \leq^{*} f$, we can only say that $f =^{*} g$, so that f and g are equal at all but finitely many values. To ease working with this relation, we weaken equality to equality modulo finite, too, and thereby make it into a partial order.

Now, we are finally ready to observe cardinal characteristic, that are fairly simple to introduce and understand. Pleasantly, they are also relatively easy to compare to other cardinalities, and because of that, we will see them somewhat frequently.

Definition 2.1. A family $\mathcal{D} \subset {}^{\omega}\omega$ is *dominating* if for each $f \in {}^{\omega}\omega$ there is $g \in \mathcal{D}$ with $f \leq^{*} g$. The *dominating number* $\mathfrak{d}$ is the smallest cardinality of any dominating family, $\mathfrak{d} = \min\{|\mathcal{D}| : \mathcal{D} \text{ dominating}\}$.

Definition 2.2. A family $\mathcal{B} \subset {}^\omega\omega$ is *unbounded* if there exists no $f \in {}^\omega\omega$ such that for all $g \in \mathcal{B}$ we have $g \leq^* f$. The *bounding number* (or perhaps more logically, *unbounding number*) $\mathfrak{b}$ is the smallest cardinality of any unbounded family, $\mathfrak{b} = \min\{|\mathcal{B}| : \mathcal{B} \text{ unbounded}\}$.

Note that $\mathfrak{b}$ would be equal to $\aleph_0$ if we used $\leq$ rather than $\leq^*$; in this case the family of all constant functions would even be unbounded.

On the other hand $\mathfrak{d}$ would be unchanged, as we could make any family $\mathcal{D}$ that was dominating under $\leq^*$, into a family that dominates under $\leq$ by adding all finite modifications of all its elements to $\mathcal{D}$. Since the number of modifications is equal to $\mathfrak{d} \times \aleph_0 \times \aleph_0$, and $\mathfrak{d} \geq \aleph_0$, we see that adding these modifications does not increase the cardinality of the family.

Now, even with the very limited set of tools we have, we can already work with these cardinalities: specifically, we can prove the following theorem.

Theorem 2.3. $\aleph_0 < \mathfrak{b} \leq \mathfrak{d} \leq \mathfrak{c}$

Proof. Let's start this proof from the left-most inequality. $\aleph_0 < \mathfrak{b}$ indicates that no countable family of functions $\mathcal{B} \subset {}^\omega\omega$ can be unbounded - in fact, we can explicitly give an upper bound for a countable family. Let $\mathcal{B}$ be a such family, and enumerate its members g_n. Then, define the function $f(x) = max_{n \leq x}(g_n(x))$. Now, for all g_n, after the first n values, f is greater than g_n, so we have $g_n \leq^* f$ for all $n \in \omega$.

Now, we will prove that $\mathfrak{b} \leq \mathfrak{d}$. Assume we have a family $\mathcal{D}$ such that the cardinality of $\mathcal{D}$ is $\mathfrak{d}$, and $\mathcal{D}$ is dominating. Assume $\mathcal{D}$ is bounded. Then there exists a single $f \in {}^\omega\omega$ $g \leq^* f$ for all $g \in \mathcal{D}$. Since we have shown before that $\leq^*$ is anti-symmetrical, this means there is no function in $\mathcal{D}$ that dominates f, which is a contradiction with our assumption that $\mathcal{D}$ is dominating. We conclude that all dominating families are also unbounded, and following this, that $\mathfrak{b} \leq \mathfrak{d}$.

Now, we have $|{}^\omega\omega| = \mathfrak{c}$, so since the largest possible family following these criteria would be the entire set ${}^\omega\omega$, it follows that $\mathfrak{c}$ is an upper bound to these cardinalities. □

Luckily, $\mathfrak{b}$ and $\mathfrak{d}$ are particularly easy to work with - even the case for equality can be specified succinctly:

Theorem 2.4. $\mathfrak{b} = \mathfrak{d}$ if and only if there is a scale in ${}^\omega\omega$, i.e. a dominating family well-ordered by $\leq *$.

Proof. Assume first that $\mathfrak{b} = \mathfrak{d}$. That means we can create a dominating family of functions $\mathcal{D}$ of size $\mathfrak{b}$, given by $\mathcal{D} = \{f_\xi : \xi < \mathfrak{b}\}$. Following this, we can create a scale by picking for each $f_\xi \in \mathcal{D}$ a g_ξ dominating f_ξ and all g_η

with $\eta < \xi$. Here we make use of the well-order of the ordinal numbers: Since we have $g_\eta < g_\xi \leftrightarrow \eta < \xi$, the family of g_i is well-ordered, because the ordinal numbers are, too. Since this family is obtained by increasing elements of a dominating family, the family itself is also dominating, and so we have found our scale.

To complete our proof, assume we have a scale, $\mathcal{C}$. Let $\mathcal{B}$ be an unbounded family of size $\mathfrak{b}$. Since $\mathcal{C}$ is dominating, we have $|\mathcal{C}| \geq \mathfrak{d} \geq \mathfrak{b}$, so we can increase every element of $\mathcal{B}$ to form a new family $\mathcal{B}' \subseteq \mathcal{C}$. Suppose that $\mathcal{B}'$ is not dominating, so there exists f such that there is no g in $\mathcal{B}'$ where $f \leq^* g$. By definition of $\mathcal{C}$, there does exist a function h such that $f \leq^* h$.

Since $\mathcal{C}$ is well-ordered, and thus totally ordered, for every $g \in \mathcal{B}'$ we have either $g \leq^* h$ or $h \leq^* g$. If $h \leq *g$, then $f \leq^* h \leq^* g$, which is in contradiction with our earlier assumption. The other option is that for every $g \in \mathcal{B}'$ we have $g \leq^* h$, which means h is an upper bound for $\mathcal{B}$, which contradicts the unboundedness of this family.

We conclude that $\mathcal{B}'$ must be a dominating family of functions of cardinality $\mathfrak{b}$, so $\mathfrak{b} = \mathfrak{d}$. □

Now, as mentioned before, these cardinal characteristics are easy to work with. Because of this, we'll be using them as a basis to work from - much of our theory will start here. Because of this, it is helpful that there are other ways of regarding these cardinalities, and the following definition shows another way:

Definition 2.5. An *interval partition* is a partition of ω into infinitely many finite intervals $I_n, n \in \omega$. These intervals are non-overlapping and numbered in the natural way: if i_n is the left endpoint of I_n for all $n \in \omega$, then $i_0 = 0, I_n = [i_n, i_{n+1})$. In this context, an interval partition $I = \{I_n : n \in \omega\}$ *dominates* another interval partition $J = \{J_n : n \in \omega\}$ if $\forall^\infty n \exists k : (J_k \subseteq I_n)$.

Theorem 2.6. $\mathfrak{d}$ is the smallest cardinality of any family of interval partitions dominating all interval partitions. $\mathfrak{b}$ is the smallest cardinality of interval partitions such that no one interval partition dominates all members of this family.

This theorem probably looks familiar - if we replace "interval partitions" with "functions", we retrieve definition 1.1 and 1.2. The proof aims to make use of this similarity, and in the following chapter we will formalize the process of doing so.

Proof. Suppose we have a dominating family $\mathcal{F}$ of interval partitions. Associate with each interval partition $I = \{I_n = [i_n, i_{n+1})$ the function $f : \omega \to \omega$, defined by letting f(x) be the right endpoint of the interval after the one containing x, so if $x \in I_n$, then $f(x) = i_{n+2} - 1$. More intuitive definitions of f will also yield the same result, but unnecessarily obfuscate the proof. This yields a family of functions, named $\mathcal{F}'$. The critical statement is now that this $\mathcal{F}'$ is a dominating family - this would prove that at least this cardinal is greater than or equal to $\mathfrak{d}$. To prove this, we have to show that for any $g \in {}^\omega\omega$ there is $f \in \mathcal{F}'$ such that $g \leq^* f$. To do this, we first have to translate g to an interval partition J. After picking an interval partition I that dominates J we can then translate I to find our desired function that dominates g.

Our way to translate back will be the following: associate to $g(x)$ an interval partition $J = \{J_n = [j_n, j_{n+1}), n \in \omega\}$, such that whenever $x \leq j_n$, then $g(x) < j_{n+1}$. A partition satisfying this condition can be given by

$$j_0 = 0, j_1 = g(1) + 1, j_n = \max_{j_{n-2} \leq i \leq j_{n-1}} (g(i)) + 1$$

Now, let I be an interval partition dominating J. Previously we associated with every interval a function $f(x)$ - let $f(x)$ be the function associated with I_n by this protocol. Then we will now show $g \leq^* f$. We will show that $g(x) \leq f(x)$ for x sufficiently large; this is sufficient since we work modulo finite. Let n be such that $x \in I_n$, and k such that $J_k \subseteq I_{n+1}$. Choosing k like this is always possible for x sufficiently large, since J is dominated by I. Now proving our desired inequality comes down to following definitions:

$$g(x) \leq j_{k+1} - 1 \leq i_{n+1} - 1 = f(x)$$

And so, we have proven that our arbitrary g is dominated by an $f \in \mathcal{F}'$, and as such, that from a dominating family of interval partitions, we can create a dominating family of functions.

The opposite construction of a dominating family of interval partitions out of a dominating family of functions is entirely equivalent: Let $\mathcal{D}$ be a dominating family of functions. Associate with each $f \in \mathcal{D}$ an interval partition as above. Suppose J is an arbitrary interval partition. Translate this to a function g as above. Then there is a function $f \in \mathcal{D}$ such that f dominates g. With this f is associated an interval partition I. Then I dominates J, with again the proof being one straightforward inequality.

The proof of the second statement can be given in a similar way, but since we're about to introduce more a more efficient method of proof, we choose not to become too repetitive, but instead, outline the proof at a later point. $\square$

3 Galois-Tukey Connections

While the previous section got us some hands-on experience with cardinal characteristics, we were working on specifics of the problem at hand that allowed us to make elegant connections. Since it might not always be clear how to form these comparisons, it is useful to provide a tool by which to formalize the idea we're using here. This part is inspired on Chapter 4 from Blass. [1] For this, we will use the following definition:

Definition 3.1. A triple $\mathbf{A} = (A_-, A_+, A)$ consisting of two sets A_- and A_+, and a binary relation $A \subset A_- \times A_+$ is, in its entirety, a *relation.* In this relation, we refer to A_- as the set of *challenges*, to A_+ as the set of *responses*, and by xAy we mean that response y *meets* challenge x.

Definition 3.2. The *norm* $\|\mathbf{A}\|$ of a relation $\mathbf{A} = (A_-, A_+, A)$ is defined as the smallest cardinality of any subset Y of A_+ such that every $x \in A_-$ is met by at least one $y \in A_+$. In other words, this is the least number of responses needed to meet all challenges.

Conveniently, the previously introduced cardinal characteristics are defined by the norm of a relation. Specifically, the following relations are relevant:

- $\mathbf{D} : ({}^\omega\omega, {}^\omega\omega,$ is dominated by$), \|\mathbf{D}\| = \mathfrak{d}$
- $\mathbf{B} : ({}^\omega\omega, {}^\omega\omega,$ does not dominate$), \|\mathbf{B}\| = \mathfrak{b}$
- $\mathbf{D}' :$ (IP, IP, is dominated by), $\|\mathbf{D}'\| = \mathfrak{d}$
- $\mathbf{B}' :$ (IP, IP, does not dominate), $\|\mathbf{B}'\| = \mathfrak{B}$

Here IP is the set of all interval partitions, and dominating is defined as earlier in this article. This shows that this new notation is certainly more efficient, but not necessarily useful - for that, we need two more concepts. Firstly, there is a connection between the cardinal characteristics as represented here, given by the notion of duality:

Definition 3.3. If $\mathbf{A} = (A_-, A_+, A)$, then the *dual* of $\mathbf{A}$ is the relation $\mathbf{A}^\perp = (A_+, A_-, A^\perp)$, where $(x, y) \in A^\perp \leftrightarrow (y, x) \notin A$.

By checking the definition, we find that the relation of being dual is symmetric. This matches our intuition of duality, which is pleasant. We can also see that from our earlier cardinal characteristics, we have 2 pairs of duals: we have $\mathfrak{D}$ and $\mathfrak{B}$ being duals, and $\mathfrak{D}'$ and $\mathfrak{B}'$ being duals.

Still, this doesn't show us any way to compare cardinalities, so the following definition will finally show us why the above helps us in this struggle:

Definition 3.4. A *morphism* from one relation $\mathbf{A} = (A_-, A_+, A)$ to another $\mathbf{B} = (B_-, B_+, B)$ is a pair of functions (ϕ_-, ϕ_+) such that

- $\phi_- : B_- \to A_-$
- $\phi_+ : A_+ \to B_+$
- $\forall b \in B_-, \forall a \in A_+$, we have if $\phi_-(b)Aa$, then $bB\phi_+(a)$.

This final condition will be known as the morphism condition. We denote a morphism by $\phi : \mathbf{A} \to \mathbf{B}$.

In this definition, we borrow the terminology "morphism" from Andreas Blass 4.8 [citation needed] rather than the original "generalized Galois-Tukey Connection" that the title speaks of which was the original name. There are several reasons for this, the primary of which is the fact that a morphism from $\mathbf{A}$ to $\mathbf{B}$ would be a generalized Galois-Tukey Connection from $\mathbf{B}$ to $\mathbf{A}$, and this direction switch might make things more intuitive. Brevity is another reason - "generalized Galois-Tukey Connection" might be more accurate, but given how it has some similarities to morphisms in other fields, it would seem "morphism" is equally fitting.

It is clear that relations and morphisms form a category R. Moreover, the operation $(A_-, A_+, A) \to (A_+, A_-, A^\perp)$, with $A^\perp$ as in definition 3.3 defines an isomorphism $R \to R^{op}$, the opposite category of R - a duality.

Theorem 3.5. If there is a morphism $\phi : \mathbf{A} \to \mathbf{B}$, then $\|\mathbf{A}\| \geq \|\mathbf{B}\|$

Proof. Let $X \subseteq A_+$ have cardinality $\|\mathbf{A}\|$ and contain responses meeting all challenges in A_-. Then $Y = \phi_+(X) \subseteq B_+$ has cardinality $\leq \|\mathbf{A}\|$, so if we prove that Y meets all challenges in B_-, that completes the first part of our proof.

Consider $b \in B_-$. By the morphism, we have $\phi_-(b) \in A_-$, and the response $x \in X$ to this challenge. Now, by the definition of a morphism, we have $\phi_-(b)Ax$, then $bB\phi_+(x)$, so $\phi_+(x)$ meets b. Since b is arbitrary, and $\phi_+(x) \in Y$, this means Y answers all challenges in B_-, and so that $\|\mathbf{B}\| \leq \|\mathbf{A}\|$. □

Theorem 3.6. If there is a morphism $\phi : \mathbf{A} \to \mathbf{B}$, then there is a morphism $\phi^\perp : \mathbf{B}^\perp \to \mathbf{A}^\perp$.

Proof. If $\phi : \mathbf{A} \to \mathbf{B}$ is a morphism given by the pair of functions (ϕ_-, ϕ_+), then $\phi^\perp : \mathbf{B}^\perp \to \mathbf{A}^\perp$ is given by (ϕ_+, ϕ_-). The proof of this statement is straightforward from the definition of a duality and a morphism. □

Finally, we now have our definitions together - we have our tool to compare cardinal characteristics, using our previously erected framework. Looking back to chapter 1, it turns out that the proofs used there all turn out to be the creation of morphisms, and the application of Theorem 2.5. To delve deeper into one case, let's look at the created morphism between $\mathcal{D}$ and $\mathcal{D}'$ as defined earlier in this chapter:

- $\mathbf{D}$: $({}^\omega\omega, {}^\omega\omega$, is dominated by)
- $\mathbf{D}'$: (IP, IP, is dominated by)

Then the function $\phi_- : \mathrm{IP} \to {}^\omega\omega$ would be the function that takes an interval partition $I_n : [i_n, i_{n+1}), n \in \omega$ to a function $f(x)$ where $x \in I_n \to f(x) = i_{n+2} - 1$. The function $\phi_+ : {}^\omega\omega \to \mathrm{IP}$ would be the function that associates with a function $f(x)$ an interval partition $I_n : [i_n, i_{n+1}), n \in \omega$ such that if $x \leq i_n$, then $f(x) < i_{n+1}$.

The remainder of the proof verified the morphism condition for both this morphism $\phi = (\phi_-, \phi_+)$ and its counterpart (ϕ_+, ϕ_-). By Theorem 2.5, this means that indeed $\mathcal{D} = \mathcal{D}'$, but also that $\mathcal{B} = \mathcal{B}'$, by duality! It is clear that this was indeed the more efficient way of proving this equality, over writing out another proof of the type given before.

A pleasant side note for morphisms is the fact that composition of morphisms works as expected: If we have a morphism $\phi : \mathbf{A} \to \mathbf{B}$, and a morphism $\psi : \mathbf{B} \to \mathbf{C}$, we can compose these to a new morphism $\chi : \mathbf{A} \to \mathbf{C}$, by letting $\chi_- = \phi_- \circ \psi_-, \chi_+ = \psi_+ \circ \phi_+$. The morphism condition follows from ψ and ϕ being morphisms.

As a closing note for this chapter, suppose for a moment we are working under CH. In that case, these cardinal characteristics are all equal to $\mathfrak{c}$. That said, one would hope the proofs as given in the previous chapter still have combinatorial value, even if the results are trivialised by CH. Sadly, this too is not the case, as shown by the following theorem:

Theorem 3.7. Let $\mathbf{A} = (A_-, A_+, A)$ and $\mathbf{B} = (B_-, B_+, B)$ be two relations, let κ be an infinite cardinal.

1. $\|\mathbf{A}\| \leq \kappa$ if and only if there is a morphism from $(\kappa, \kappa, =)$ to $\mathbf{A}$.
2. If $\|\mathbf{A}\| = |A_+| = \kappa$, then there is a morphism from $\mathbf{A}$ to $(\kappa, \kappa, \leq)$.
3. If $\|\mathbf{A}^\perp\| = |A_-| = \kappa$, then there is a morphism from $(\kappa, \kappa, \leq)$ to $\mathbf{A}$.
4. If $\|\mathbf{A}| = |A_+| = \|\mathbf{B}^\perp\| = |B_-| = \kappa$, then there is a morphism from $\mathbf{A}$ to $\mathbf{B}$.

This last point is our problem: it says that morphisms in both directions exist between relations with equal norm, and this means that, after choosing a model, morphisms are only interesting if their results are.

Proof. Part 1: If there is a morphism from $(\kappa, \kappa, =)$ to $\mathbf{A}$, then $\|\mathbf{A}\| \leq \kappa$ follows from a previous theorem. The other direction is more difficult: Assume $\|\mathbf{A}\| \leq \kappa$. We want to create a pair of functions (ϕ_-, ϕ_+) such that the morphism condition holds, which in this case would be

$$\phi_-(a) = k \to aA\phi_+(k)$$

To do this, let $\phi_+ : \kappa \to A_+$list the responses to all challenges in A_-. By our assumption, there are no more than κ responses necessary to do this, so this construction can be completed. Since we have listed all necessary responses, we can define $\phi_-(a)$ to be α so that $\phi_+(\alpha)$ meets a. After defining $\phi = (\phi_i, \phi_+)$ in this manner, checking the morphism condition is straightforward from the definition, and we conclude that as such, a morphism exists from $(\kappa, \kappa, =)$ to $\mathbf{A}$.

Part 2: Assume $\|\mathbf{A}^\perp\| = |A_-| = \kappa$. Now, let $\phi_+ : A_+ \to \kappa$ be any injective function. For any $\alpha < \kappa$ the set $\{a \in A_+ : \phi_+(a) \leq \alpha\}$ has cardinality smaller than κ since ϕ_+ is injective, so there is some challenge in A_- that is unmet. Let $\phi_-(\alpha)$ be any such challenge.

Now, we have to check the morphism condition:

$$\forall k \in \kappa, \forall a \in A_- : \phi_-(k)Aa \to k < \phi_+(A)$$

Fix $k \in \kappa$. Then, there is an $a \in A_-$ such that a meets $\phi_-(k)$. Now, by our choice of ϕ_-, this a cannot be an element of the set $\{a \in A_+ : \phi_+(a) \leq k\}$. Then, necessarily, $k < \phi_+(A)$, so our condition is met.

Part 3: Assume $\|\mathbf{A}^\perp\| = |A_-| = \kappa$. Then we can create a morphism from $\mathbf{A}^\perp$ to $(\kappa, \kappa, \leq)$, by part 2. Taking the dual morphism of this one gives the desired morphism from $(\kappa, \kappa, \leq)$ to $\mathbf{A}$.

Part 4: From part 2, we can create a morphism from $\mathbf{A}$ to $(\kappa, \kappa, \leq)$, from part 3 there is a morphism from $(\kappa, \kappa, \leq)$ to $\mathbf{B}$, and we had established morphisms can be composed. $\square$

4 Further cardinal characteristics

Now that we have all relevant tools in place, let us look at some further cardinal characteristics, taken from Blass' chapter 3 and 5. [1]

Definition 4.1. A set $X \subseteq \omega$ splits an infinite set $Y \subseteq \omega$ if both $Y \cap X$ and $Y - X$ are infinite. A splitting family $\mathcal{S}$ of subsets of ω such that each infinite $Y \subseteq \omega$ is split by at least one $X \in S$. The splitting number $\mathfrak{s}$ is the smallest cardinality of any splitting family.

Rephrasing in the notation of the past chapter, $\mathbf{S} = (\mathcal{P}(\omega), \mathcal{P}(\omega)$, is split by).

Theorem 4.2. $\mathfrak{s} \leq \mathfrak{d}$.

Proof. Let $\mathbf{D}'$ =(IP, IP, is dominated by). Now, since we know $\|\mathbf{D}'\|$ we want to create a morphism from $\mathbf{D}'$ to $\mathbf{S}$. Let $\phi_- : \mathcal{P}(\omega) \to \mathrm{IP}$ be the function that takes a set $X \subset \omega$ to an interval partition $I = I_n : n \in \omega$ such that for all n, there is at least one $x \in X$ contained in I_n.

Take ϕ_+ :IP $\to \mathcal{P}(\omega)$, that takes an interval partition $I = I_n : n \in \omega$ to the subset $\cup_{n\in\omega} I_{2n}$, so every even interval in the partition.

Now, the only thing to check is the morphism condition, which would, in this case, be phrased as

$$\forall s \in \mathcal{P}(\omega), \forall d \in \mathrm{IP} : \phi_-(s) \text{ is dominated by } d \to s \text{ is split by } \phi_+(d)$$

To prove this, fix an interval partition s. By definition, there is a $d \in \mathrm{IP}$ such that $\phi_-(s)$ is dominated by d. Then, by the definition of interval partitions dominating, we have (modulo finite) that every interval of d contains an interval of $\phi_-(s)$. Since every interval of $\phi_-(s)$ contains an element of the set s, so does d. Now, since $\phi_+(d)$ is the union of even-numbered intervals, and each intervals contains a member of s, it's clear $\phi_+(d) \cap s$ is infinite. Since $\phi_+(d) - s$ is the union of odd-numbered intervals, by the same argument, this set, too, is infinite.

We conclude that this is indeed a morphism, and so, that $\mathfrak{s} \leq \mathfrak{d}$. □

Apart from these cardinal invariants that have a set definition, we can also define standard characteristics, that can be applied to any ideal of a set.

Definition 4.3. An *ideal* of subsets of a set X is a family $\mathcal{I}$ of subsets which is downwards closed (if $I \in \mathcal{I}$ and $J \subseteq I$ then $J \in \mathcal{I}$) and closed under finite unions ($\emptyset \in \mathcal{I}$ and we have $\cup_n I_n \in \mathcal{I}$ for finite collections of I_n). In this article, we also assume all singletons of X to be in I, so if $x \in X$, then $\{x\} \in \mathcal{I}$. A σ-ideal is an ideal that's closed under countable unions.

On an ideal, we can then define the following cardinal invariants:

Definition 4.4. Let $\mathcal{I}$ be a proper ideal of a set X. The on this ideal $\mathcal{I}$, we can define

- $\mathbf{add}(\mathcal{I})$, the additivity of $\mathcal{I}$, the smallest number of sets in $\mathcal{I}$ with union not in $\mathcal{I}$.
- $\mathbf{cov}(\mathcal{I})$, the covering number of $\mathcal{I}$, the smallest number of sets in $\mathcal{I}$ with union X.
- $\mathbf{non}(\mathcal{I})$, the smallest cardinality of any subset of X not in $\mathcal{I}$.
- $\mathbf{cof}(\mathcal{I})$, the smallest cardinality of any subset $\mathcal{B}$ of $\mathcal{I}$ such that every element of $\mathcal{I}$ is a subset of an element of $\mathcal{B}$. In this case, $\mathcal{B}$ is a basis for $\mathcal{I}$.

However, these definitions do not make use of our previously set up framework, and that is something we can change - it turns out they fit in very neatly. We can rephrase the above definition by defining the following two relations:
$\mathbf{Cov}(\mathcal{I}) = (X, \mathcal{I}, \in)$
$\mathbf{Cof}(\mathcal{I}) = (\mathcal{I}, \mathcal{I}, \subseteq)$
Then, we have the following connections: $\|\mathbf{Cov}(\mathcal{I})\|$ can be phrased as "the smallest number of sets in $\mathcal{I}$ such that every element in X is in one of these sets", which means these sets would have union X. That means $\|\mathbf{Cov}(\mathcal{I})\| = \mathbf{cov}(\mathcal{I})$ - it is always pleasant when a choice of name turns out to have an intuitive result.

Now, $\mathbf{Cov}(\mathcal{I})^\perp = (\mathcal{I}, X, \not\ni)$. This can be worded as "the smallest number of elements in X such that no set in $\mathcal{I}$ contains all these elements." Here, we can clearly see why we don't allow $X \in \mathcal{I}$, as that would create difficulties here. Now, suppose A is a set of elements satisfying this condition, so for every set $I \in \mathcal{I}$ there is an element $a \in A$ such that $a \notin I$.

Clearly, this set A cannot be in $\mathcal{I}$ itself, as there is no element in A that is not in A. We also propose that every subset of X that is not a member of $\mathcal{I}$ has the property that its difference with every set in $\mathcal{I}$ is non-empty.

Taking these two statements together, we see that this set is equal to the smallest subset of X not in $\mathcal{I}$, or $\|\mathbf{Cov}(\mathcal{I})^\perp\| = \mathbf{non}(\mathcal{I})$.

$\|\mathbf{Cof}(\mathcal{I})\|$ can be phrased as "The smallest number of sets in $\mathcal{I}$ such that every set in $\mathcal{I}$ is a subset of one of these sets", so $\|\mathbf{Cof}(\mathcal{I})\| = \mathbf{cof}(\mathcal{I})$.

The approach for $\|\mathbf{Cof}(\mathcal{I})^\perp\|$ is very similar - $\mathbf{Cof}(\mathcal{I})^\perp = (\mathcal{I}, \mathcal{I}, \not\supseteq)$. In an analagous way to $\mathbf{Cov}(\mathcal{I})^\perp$, we find that this is equal to $\mathbf{add}(\mathcal{I})$.

Now that we have this in place, we can start to prove inequalities the way we had established before: Morphisms. There is one simple morphism, with a free dual result:

Theorem 4.5. $\mathbf{cof}(\mathcal{I}) \geq \mathbf{cov}(\mathcal{I})$.

Proof. We want to create a morphism from $\mathbf{Cof}(\mathcal{I}) = (\mathcal{I}, \mathcal{I}, \subseteq)$ to $\mathbf{Cov}(\mathcal{I}) = (X, \mathcal{I}, \in)$, so we require a function $\phi_- : X \to \mathcal{I}$, and a function $\phi_+ : \mathcal{I} \to \mathcal{I}$ such that we have

If $\phi_-(x) \subseteq I$, then $x \in \phi_+(I)$ for all $x \in X$, $I \in \mathcal{I}$.

It turns out these maps can be fairly simple. If we take ϕ_+ to be the identity map, and $\phi_-(x)$ to be the set containing x, the condition is satisfied: If $x \subseteq I$, then $x \in I$. We conclude this choice of ϕ creates a morphism between $\mathbf{Cov}(\mathcal{I})$ and $\mathbf{Cof}(\mathcal{I})$, and as such, that $\mathbf{cof}(\mathcal{I}) \geq \mathbf{cov}(\mathcal{I})$. □

Corollary 4.6. $\mathbf{non}(\mathcal{I}) \geq \mathbf{add}(\mathcal{I})$

There is another pair of inequalities between these cardinal invariants, but there creating an explicit morphism is much more complicated than the "informal" proof, so we will provide that verison instead.

Theorem 4.7. $\mathbf{cov}(\mathcal{I}) \geq \mathbf{add}(\mathcal{I})$

Proof. Let $\mathcal{F}$ be a family of sets in $\mathcal{I}$ with union X, and $|\mathcal{F}| = \mathbf{cov}(\mathcal{I})$. Clearly, since the union of all sets in $\mathcal{F}$ is X, this is outside $\mathcal{I}$, and so, $\|\mathcal{F}\| \geq \mathbf{add}(\mathcal{I})$. We conclude that the theorem holds. □

Now, even though we have not explicitly given a morphism, we can still get the dual result. According to Theorem 3.7, there exists a morphism between any two comparable cardinal invariants. Specifically, there is a morphism between $\mathbf{Cov}(\mathcal{I})$ and $\mathbf{Cof}(\mathcal{I})^{\perp}$.

This means that the dual morphism exists, too, and as such that we do still have the dual result $\mathbf{cof}(\mathcal{I}) \geq \mathbf{non}(\mathcal{I})$.

These are all the commonly used cardinal invariants on ideals, so now that we've established some relations between them, let's take a look at their cardinality. In this thesis we'll only be working with σ-ideals, so clearly $\mathbf{add}(\mathcal{I})$ is uncountable, and so all other invariants are, too. We choose the underlying set X to be the continuum, so we have $\mathbf{cof}(\mathcal{I}) \leq \mathfrak{c}$, and so the others are too, and we find that at least these cardinal invariants are in the same scope as the cardinal characteristics defined earlier.

An interesting first choice for $\mathcal{I}$ is to choose $\mathcal{I} = \mathcal{K}_\sigma$, the ideal generated by the compact subsets of ${}^\omega\omega$, or the ideal of all sets coverable by countably many compacts.

Theorem 4.8. (Blass [1] 2.8) $\mathbf{add}(\mathcal{K}_\sigma) = \mathbf{non}(\mathcal{K}_\sigma) = \mathfrak{b}$ and $\mathbf{cov}(\mathcal{K}_\sigma) = \mathbf{cof}(\mathcal{K}_\sigma) = \mathfrak{d}$.

Proof. Since a subset of a discrete space ω is compact iff it is finite, the Tychonoff theorem implies that a subset of ${}^\omega\omega$ is compact iff it is closed and included in a product of finite subsets of ω. There is no loss of generality in taking the finite subsets to be initial segments, so we find that all sets of the form

$$\{f \in {}^\omega\omega : f \leq g\} = \prod_{n\in\omega}[0, g(n)]$$

are compact, and all compacts are of this form. It follows that every set of this form is in $\mathcal{K}_\sigma$, and every set in $\mathcal{K}_\sigma$ is a subset of one of these. We have now created an equivalence between these sets, and so a morphism can be formed by this mapping, which completes the theorem. □

So, in an unexpected place, we again find our initial cardinal characteristics, $\mathfrak{b}$ and $\mathfrak{d}$, yielding yet another way to work with these if the occasion calls for it.

However, there are also choices for $\mathcal{I}$ that are not generally equal to other cardinal characteristics, and as such, that are interesting to investigate by themselves. There are 2 ideals which are the most prevalent choices to investigate: There is $\mathcal{B}$, the ideal of meager sets (after Baire), and $\mathcal{L}$, the ideal of sets of Lebesgue measure zero (after Lebesgue).

Before we can work with these, let's reassure ourselves of these definitions:

Definition 4.9. Let X be a complete seperable metric space. A set $A \subset X$ is nowhere dense if the complement of A contains a dense open set. A set $A \subset X$ is meager if A is the union of countably many nowhere dense sets. A set $A \subset X$ is comeager if the complement of A is meager, or if it is the intersection of countably many dense sets.

These definitions are good to know, but rather impractical to work with, so before looking at our cardinal invariants, we will start out by giving a more convenient description of meagerness in ${}^\omega 2$.

Definition 4.10. A chopped real is a pair (x, Π), where $x \in {}^\omega 2, \Pi \in \mathrm{IP}$. We write $\mathrm{CR} = {}^\omega 2 \times \mathrm{IP}$ for the set of chopped reals. A real $y \in {}^\omega 2$ *matches* a chopped real (x, Π) if $x|_I = y|_I$ for infinitely many intervals $I \in \Pi$.

Theorem 4.11. A subset M of ${}^\omega 2$ is meager iff there is a chopped real that no member of M matches.

Proof. The set of reals that match a given chopped real $(x, \{I_n : n \in \omega\})$ is

$$\bigcap_k \bigcup_{n \geq k} \{y : x|_{I_n} = y|_{I_n}\}$$

Each of these unions is dense, there are countably many unions, and as such, this intersection is comeager. This means all reals that do not match a particular real form a meager set. Any subset M of ${}^\omega 2$ that has a chopped real that no member of M matches is then a subset of a meager set, and in that, itself meager.

The proof of the other direction goes beyond the scope of this thesis but can be found in Blass [CITATION NEEDED]. $\square$

Now, we have meagerness defined in interval partitions and reals, we can use this to compare these cardinal invariants to the ones related to IP, which again turn out to be $\mathfrak{b}$ and $\mathfrak{d}$. To do so, we need the following characterisation:

Proposition 4.12. Let $\mathrm{Match}(x, \Pi)$ be the set of reals matching the chopped real (x, Π). Then, we have $\mathrm{Match}(x, \Pi) \subseteq \mathrm{Match}(x', \Pi')$ if and only iff for all but finitely many intervals $I \in \Pi$ there exists an interval $J \in \Pi'$ such that $J \subseteq I$ and $x'|_J = x|_J$. In this case, we say (x, Π) engulfs (x', Π').

The proof of this statement is straightforward, and as such, omitted. The outline of the proof is that a real matches (x, Π) if it agrees with x on infinitely many intervals in Π, and since each of these intervals contain an interval in Π' where x and x' agree, that means every real that matches (x, Π) also matches (x', Π').

Now, we can rephrase $\mathbf{Cof}(\mathcal{B})$, which was previously stated as $(\mathcal{B}, \mathcal{B}, \subseteq)$, to $\mathbf{Cof'}(\mathcal{B})$ = (CR, CR, is engulfed by).

Note also that if (x, Π) engulfs (x', Π'), then Π dominates Π'. So, there is a trivial morphism from $\mathbf{Cof'}(\mathcal{B})$ to $\mathcal{D}'$ as defined earlier. Recall that CR $= {}^\omega 2 \times$ IP, so this morphism is very easy to form.

This gives us the following inequalities:

Corollary 4.13. $\mathbf{add}(\mathcal{B}) \leq \mathfrak{b}$ and $\mathfrak{d} \leq \mathbf{cof}(\mathcal{B})$

Similarly, we can rephrase $\mathbf{Cov}(\mathcal{B})$ to $\mathbf{Cov}'(\mathcal{B})$ = (${}^\omega 2$, CR, does not match). Now, we can use our previous knowledge of $\mathcal{K}_\sigma$ to prove the following theorem:

Theorem 4.14. $\mathbf{cov}(\mathcal{B}) \leq \mathfrak{d}$ and $\mathfrak{b} \leq \mathbf{non}(\mathcal{B})$

Proof. Firstly, we assert that $\mathcal{K}_\sigma \subseteq \mathcal{B}$. Recall from the start of this section that the compact subsets of ${}^\omega\omega$ are given by $f : f \leq g$, and these sets are nowhere dense. It is fairly straightforward to see that if $\mathcal{I} \subseteq \mathcal{J}$ are ideals, that $\mathbf{cov}(\mathcal{J}) \leq \mathbf{cov}(\mathcal{I})$, since any family of subsets with union X in $\mathcal{I}$ are also in $\mathcal{J}$. Similarly, the inequality $\mathbf{cof}(\mathcal{J}) \leq \mathbf{cof}(\mathcal{I})$ holds for similar reasons.

We can conclude that $\mathbf{cov}(\mathcal{B}) \leq \mathbf{cov}(\mathcal{K}_\sigma)$, or $\mathbf{cov}(\mathcal{B}) \leq \mathfrak{d}$. From duality, the second part of the theorem immediately follows. □

This concludes all relevant inequalities on the ideal of meager sets, but before we move on to the sets of measure zero, it's convenient to look at an alternative, combinatorial description of $\mathbf{cov}(\mathcal{B})$:

Definition 4.15. Call two functions *eventually different* if $\forall^\infty n(x(n) \neq y(n))$. Otherwise, call them *infinitely equal*, that is, if $\exists^\infty n(x(n) = y(n))$.

Theorem 4.16. $\mathbf{cov}(\mathcal{B}) = \|({}^\omega\omega, {}^\omega\omega, \text{eventually different})\|$

Proof. In this thesis we will only prove $\mathbf{Cov}(\mathcal{B}) \leq \|({}^\omega\omega, {}^\omega\omega, \text{eventually different})\|$; the opposite inequality uses theory that is not relevant for any further part of this thesis, but can be found in Blass [CITATION NEEDED] we will construct a morphism ϕ. For this, recall first that $\mathbf{Cov}(\mathcal{B}) = ({}^\omega\omega, \mathcal{B}, \in)$.

Define $\phi_- : {}^\omega\omega \to {}^\omega\omega$ to be the identity map, and $\phi_+ : {}^\omega\omega \to \mathcal{B}$ the map that sends an element $a \in {}^\omega\omega$ to the set of $b \in {}^\omega\omega$ such that b is eventually different from a.

The related morphism condition then becomes: For all $x \in {}^\omega\omega$, for all $y \in {}^\omega\omega$, if x is eventually different from y, then x is an element of the set of elements of ${}^\omega\omega$ eventually different from y, which is evidently true.

The only thing that remains to check now is the fact that, given fixed x, the set of all y eventually different from x is meager.

The set of reals that are infinitely equal to a given $x \in \omega\omega$ is

$$\bigcap_k \bigcup_{n \geq k} \{y : x(n) = y(n)\}$$

A construction like this has been seen before in theorem 4.11, and as was the case there, this set is comeager. Its complement is the set we were looking for, and consequently, this set is meager. □

Since we have created a morphism, we get the dual result for free:

Corollary 4.17. $\mathbf{non}(\mathcal{B}) = \|({}^\omega\omega, {}^\omega\omega, \text{infinitely equal})\|$

Since these definitions don't inherently depend on the definition of meager sets, they are much easier to compare to cardinal invariants with different underlying ideal. Specifically, the other ideal that is interesting is the σ-ideal of sets of measure zero. Sadly, there is no direct causality between sets being meager and their measure being zero, so we need to do some work here. The following theorem shows part of this work:

Theorem 4.18. $\mathbf{cov}(\mathcal{B}) \leq \mathbf{non}(\mathcal{L})$

Proof. Let Π be the interval partition whose n^{th} interval has cardinality $n+1$ for all n. Define a binary relation R on ${}^{\omega}2$ by letting xRy if and only if there exist infinitely many n such that $x|_{I_n} = y|_{I_n}$, i.e. y matches (x, Π). We can see that R is symmetric and reflexive - transitivity does not hold, though.

Define $R_x = y : xRy$. Again, similarly to theorem 4.11, we see that R_x is comeager. We also find that R_x has measure zero. For fixed x, all y that agree with x on I_n form a set of measure $2^{-(n+1)}$ (recall that $\|I_n\| = n+1$). All y that agree with x on a single I_n with $n > k$ form a set of measure at most 2^{-k}, so seeing how all $y \in R_x$ agree with a single I_n with $n > k$ for all $k \in \omega$, we find that the measure of these y has to be 0.

Define the relation $\mathbf{R} = ({}^{\omega}2, {}^{\omega}2, R)$. We can now create a morphism $\phi : \mathbf{R} \to \mathbf{Cov}(\mathcal{L})$ by letting $\phi_- : {}^{\omega}2 \to {}^{\omega}2$ be the identity map, and $\phi_+ : {}^{\omega}2 \to \mathcal{L}$ defined by $\phi_+(x) = R_x$.

The associated morphism condition is then: For all $x \in {}^{\omega}2$, for all $y \in {}^{\omega}2$, if xRy, then $x \in R_y$, and by the symmetry of R, this is true.

Then, we can create a different morphism $\psi : \mathbf{Cov}(\mathcal{B}) \to \mathbf{R}^{\perp}$. Again, let $\psi_- : {}^{\omega}2 \to {}^{\omega}2$ be the identity map, but this time $\phi_+ : {}^{\omega}2 \to \mathcal{B}$ is defined by $\phi_+(x) = {}^{\omega}2 - R_x$. Note that since R_x is comeager, this set is indeed meager.

The associated morphism condition is then: For all $x \in {}^{\omega}2$, for all $y \in {}^{\omega}2$, if not xRy, then $y \in {}^{\omega}2 - R_x$, and again, this is clearly true.

From these morphisms, we can derive the inequalities $\mathbf{cov}(\mathcal{L}) \leq \|\mathbf{R}\|$ and $\mathbf{cov}(\mathcal{B}) \leq \|\mathbf{R}^{\perp}\|$. Then, we know the dual morphisms also exist, and because of that, we also have $\|\mathbf{R}^{\perp}\| \leq \mathbf{non}(\mathcal{L})$ and $\|\mathbf{R}\| \leq \mathbf{non}(\mathcal{B})$. We can then compose these morphisms, or at very least, these inequalities to get the inequality $\mathbf{cov}(\mathcal{B}) \leq \mathbf{non}(\mathcal{L})$ and its dual result $\mathbf{cov}(\mathcal{L}) \leq \mathbf{non}(\mathcal{B})$ □

The following theorem is given without proof, as the proof would take us through a several-page-long trip through otherwise unnecessary territory - it is relatively well documented in for example [CITATION NEEDED].

Theorem 4.19. $\mathbf{add}(\mathcal{L}) \leq \mathbf{add}(\mathcal{B})$ and $\mathbf{cof}(\mathcal{B}) \leq \mathbf{cof}(\mathcal{L})$

All inequalities given in this section are usually summarized in Cichoń's diagram as pictured below, taken from Blass [1], though originally given by Bartoszyński et al. [2]

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

This diagram represents all possible inequalities between these cardinal invariant, in the sense that an assignment of $\aleph_1$ and $\mathfrak{c}$ is consistent with this diagram if and only if it is consistent with ZFC. As such, for any axiom added to ZFC, checking the effects on the cardinal invariants in this diagram is usually informative about the nature of the resulting model.

5 Martin's Axiom

We have established a lot of inequalities over the past chapters, but no actual model to evaluate any of the presented value. The only axiom we have seen so far that assigns a value to these cardinal characteristics is the Axiom of Choice, which trivially makes all these cardinal characteristics equal to $\mathfrak{c}$ - not a very satisfying result of all the effort spent.

Another option for an axiom that allows us to work with these cardinal characteristics without having to delve too deep into model theory is Martin's Axiom. Our treatment of this axiom follows Blass' chapter 7. [1] The underlying theory is very complicated, but this is all neatly swept under the rug by working with only this axiom. To work with the axiom, we need to introduce some further new concepts, resulting in the following two definitions:

Definition 5.1. Let $(P, \leq)$ be a nonempty partially ordered set. Two elements $p, q \in P$ are compatible if they have a common lower bound, and incompatible otherwise. An antichain is a set of pairwise incompatible elements. P satisfies the countable chain condition (ccc, or also countable antichain condition) if all its antichains are countable, or more generally, P satisfies the $< \kappa$ chain condition if all its antichains have cardinalities $< \kappa$.

Definition 5.2. Let $(P, \leq)$ be a nonempty partially ordered set. A subset $D \subseteq P$ is dense if every element of P is $\geq$ an element of D. If $\mathcal{D}$ is a family of dense subsets of P, then $G \subseteq P$ is $\mathcal{D}$-generic if it is closed upward (?), every two of its members have a common lower bound, and it intersects every $D \in \mathcal{D}$. A family of subsets is a filter if it contains the empty set, is closed under intersection, and supersets.

Now, bear in mind, this is the smallest framework we can provide to understand the following axiom, so ensure a decent understanding of these concepts before continuing.

Definition 5.3. Martin's Axiom is the statement that, if $\mathcal{D}$ is a family of fewer than $\mathfrak{c}$ dense subsets of a partial order P with ccc, then there is a $\mathcal{D}$-generic filter $G \subseteq P$.

There are several generalizations possible of this notion, but to not overextend into the wild forest of forcing axioms, we will limit ourselves to only the simplest version: Martin's Axiom

Disappointingly, Martin's Axiom has the same result as AC, as proven by the following theorem:

Theorem 5.4. MA implies $\mathbf{add}(\mathcal{L}) = \mathfrak{c}$.

Proof. Let κ be a cardinal number with $\kappa < \mathfrak{c}$, and we are given κ sets $N_\alpha \subseteq \mathbb{R}$ of measure 0. If, assuming Martin's Axiom, the union of these κ sets have union with measure 0, we have proven the theorem. To do this, we will prove that for any positive ϵ there exists a set with measure ϵ with all N_α as subsets - the theorem then immediately follows.

Given ϵ, let P be the set of open subsets of $\mathbb{R}$ having measure smaller than ϵ, and order P by reverse inclusion. Before we can apply Martin's Axiom, we first need to verify this set satisfies the ccc.

To do this, let uncountably many elements p of P be given. Find inside of all these open subsets a finite union $q(p)$ of open intervals with rational endpoints, so that $\mu(p - q(p)) < \epsilon - q(p)$; this is possible since the rationals are dense in the reals. From the definition of a measure then follows that $\mu(p - q(p)) < \frac{1}{2}(\epsilon - \mu(q(p)))$. Since the intervals have rational endpoints, there are only countably many different possibilities for $q(p)$, so there are two p with the same $q(p)$. Since their union is $q(p)$ combined with the two remainders $p - q$, each of which has norm smaller than $\frac{1}{2}(\epsilon - \mu(q(p)))$, their union has norm smaller than ϵ.

This means these two p share a common lower bound in P, and as such, there cannot be an uncountable antichain in P, so the ccc is satisfied.

Now, we can work towards applying Martin's Axiom: For each of the given N_α, let $D_\alpha = \{p \in P : N_\alpha \subseteq p\}$. Note that this is a dense subset, due to N_α having measure zero. We also have that $\kappa < \mathfrak{c}$, so finally, all conditions for Martin's Axiom are met. It then supplies us with a generic G meeting all the D_α. Then, $\bigcup G$ includes all the N_α, and due to the definition of P, we have that G also has measure smaller than ϵ, which proves the theorem. $\square$

Following Cichon's diagram, we see that this means that under Martin's Axiom, all cardinal characteristics in the diagram are equal to $\mathfrak{c}$. So, since Martin's Axiom is given as a weakening of the continuum hypothesis, we find through this diagnostic tool that the hypothesis hasn't weakened that much.

References

[1] Andreas Blass. *Combinatorial Cardinal Characteristics of the Continuum.* University of Michigan, 2003

[2] Tomek Bartoszyński, Haim Judah and Saharon Shelah. *The Cichoń Diagram.* The Journal of Symbolic Logic, Volume 58, Number 2, page 401-423, 1993