

1. Overview

We will prove consistency and independence results. These are theorems which assert that within some given axiomatic system, some given proposition cannot be disproven or can be neither proven nor disproven.


Example: the parallel postulate in Euclidean geometry states:

Given a line and a point not on the line, only one line can be drawn through the point parallel to the line.

This statement is independent of the other four axioms:

1. A straight line can be drawn between any two points.
2. A finite line can be extended infinitely in both directions.
3. A circle can be drawn with any center and any radius.
4. All right angles are equal to each other.

How do we know this? The parallel postulate is consistent with the other four axioms because all five statements hold in the standard Euclidean plane. Its negation is consistent with the other four axioms because the sphere model ("point" is a point on the sphere, "line" is a great circle) satisfies the first four Euclidean axioms, but does not satisfy the parallel postulate. Thus, the parallel postulate is independent of the other axioms. (Another model where the parallel postulate fails is the Poincaré disk.)



"Lines" in the Poincaré disk

Do the sphere and disk models really satisfy the first four Euclidean axioms? This is not perfectly clear because of the informal manner in which we stated the axioms. There seem to be some questions of interpretation, especially in the case of the sphere model. This illustrates the need to have a formal language in which our axioms can in principle be rigorously expressed.

Peano arithmetic is a good example of a rigorous, formal axiomatic system. We have an infinite list of variables x, y, \dots ; a constant symbol 0 ; symbols for addition, multiplication, and successor $(+, \cdot, ')$; and parentheses. A term is any grammatical expression built up from these components. An atomic formula is a statement of the form $t_1 = t_2$ where t_1 and t_2 are terms. A formula is any statement built up from atomic

formulas using the logical symbols \rightarrow (implies), \neg (not), and \forall (for all). (All of the other logical symbols can be defined in terms of these.)

The axioms of Peano arithmetic (PA) consist of the logical axioms:

$$\begin{aligned} &\phi \rightarrow (\psi \rightarrow \phi) \\ &[\phi \rightarrow (\psi \rightarrow \omega)] \rightarrow [(\phi \rightarrow \psi) \rightarrow (\phi \rightarrow \omega)] \\ &(\neg\psi \rightarrow \neg\phi) \rightarrow [(\neg\psi \rightarrow \phi) \rightarrow \psi] \\ &(\forall x)\phi(x) \rightarrow \phi(t) \\ &(\forall x)(\phi \rightarrow \psi) \rightarrow [\phi \rightarrow (\forall x)\psi] \end{aligned}$$

and the non-logical axioms:

$$\begin{aligned} &(x = y) \rightarrow [(x = z) \rightarrow (y = z)] \\ &(x = y) \rightarrow x' = y' \\ &\neg(0 = x') \\ &(x' = y') \rightarrow (x = y) \\ &x + 0 = x \\ &x + y' = (x + y)' \\ &x \cdot 0 = 0 \\ &x \cdot (y') = (x \cdot y) + x \\ &\phi(0) \rightarrow [(\forall x)(\phi(x) \rightarrow \phi(x')) \rightarrow (\forall x)\phi(x)]. \end{aligned}$$

These axioms are presented schematically, meaning that ϕ , ψ , and ω stand for any formulas, x , y , and z can be replaced by any variables, and t stands for any term. (There are some minor syntactic restrictions on x and t in the last two logical axioms.) In other words there are actually infinitely many axioms, each one of which fits one of the formats given above.

The formal system is completed by specifying two rules of inference:

$$\begin{aligned} &\text{from } \phi \text{ and } \phi \rightarrow \psi, \text{ infer } \psi \\ &\text{from } \phi, \text{ infer } (\forall x)\phi \end{aligned}$$

which are also presented schematically. A theorem of PA is any formula which is either an axiom or follows from other theorems by a rule of inference.

Peano arithmetic attains a perfect degree of formality. It would not be difficult to write a computer program which would mechanically list all theorems of the system. The question of whether some statement is or is not a theorem is completely precise and might even be characterized as a combinatorial question.

In the twentieth century mathematicians decided that set theory is the fundamental subject in terms of which all ordinary mathematics should in principle be interpreted. [Editorial comment: This may have been a mistake. Philosophically, the concept of a set was never very convincing and seemed to arise out of elementary grammatical confusions. (Halmos: "A pack of wolves, a bunch of grapes, or a flock of pigeons are all examples of sets of things." Black: "It ought then to make sense, at least sometimes, to speak of being pursued by a set, or eating a set, or putting a set to flight.") Moreover, proof theorists now know that all or virtually all mainstream mathematics can be formalized in essentially number-theoretic systems that are metaphysically unexceptionable. However, even if set theory does not really belong in the foundational role in which it is currently cast, it is still a beautiful and fascinating subject worthy of study in its own right.]

The formal language of set theory is even simpler than the language of Peano arithmetic. The only atomic formulas are $x \in y$ and $x = y$ (for any variables x and y). Arbitrary formulas are built up from the atomic ones using, say, the symbols \rightarrow , \neg , and \forall just as in Peano arithmetic. The rules of inference and the logical axioms are also the same as those of PA, except that now variables are the only terms. The standard non-logical axioms of set theory — the Zermelo-Fraenkel axioms including the axiom of choice (ZFC) — are not quite so neat, however. Several of them become fairly unreadable when expressed formally. We will state them informally, but it should be clear that formalizing them would be routine.

Zermelo-Fraenkel axioms

1. Extensionality. For all x and y , $x = y$ if and only if they have the same elements.
2. Pairing. For all x and y there exists a set $\{x, y\}$ whose elements are exactly x and y .
3. Separation scheme. For all x there exists a set $\{u \in x : \phi(u)\}$ whose elements are exactly those sets in x which satisfy ϕ . (One axiom for each formula ϕ .)
4. Union. For all x there exists a set $y = \bigcup_{u \in x} u$, the union of all $u \in x$.
5. Power set. For all x there exists a set $y = \mathcal{P}(x)$, the set of all subsets of x .
6. Infinity. There exists an infinite set.
7. Replacement scheme. If for every u there is at most one v such that $\phi(u, v)$, then for every x there exists a set y whose elements are exactly those v such that $\phi(u, v)$ for some $u \in x$. (One axiom for each formula ϕ .)
8. Foundation. Every nonempty set has an \in -minimal element.
9. Choice. Every set of nonempty sets has a choice function.

We also need an axiom about equality which states that for all x, y , and z , if $x = y$ and $x \in z$ then $y \in z$.

(Technical point: in separation and replacement the formula ϕ could have free variables other than those indicated. If ϕ has other free variables x_1, \dots, x_n then the axioms should be understood as asserting that for all x_1, \dots, x_n the stated assertion holds.)

Our goal is to prove that various statements are consistent with or independent of ZFC. The example of the parallel postulate mentioned above illustrates the basic idea: models. We want to show that the continuum hypothesis (CH) is consistent with ZFC by constructing a model of ZFC + CH — a structure in which both the ZFC axioms and the continuum hypothesis hold — and that it is independent of ZFC by also constructing a model of ZFC + \neg CH. This is exactly how we show that the parallel postulate and its negation are both consistent with the other Euclidean axioms. The obvious problem in the case of set theory is that it is not clear that there are any models of ZFC at all. We get around this difficulty by assuming that ZFC has models. We then use the technique of forcing to convert a given model of ZFC into a model of ZFC plus some other statement ϕ . The result is a theorem which states “if ZFC is consistent, then so is ZFC + ϕ .” We say that ϕ is relatively consistent with ZFC.

~~It is easy to get confused about where these arguments take place. A model is a set-theoretic object, so are we reasoning in ZFC the whole time? Is it legitimate to reason about the existence of models of ZFC within ZFC? If we are working in ZFC, then aren't we taking the consistency of ZFC as given? There are various ways of answering these concerns. We will use the following device. We define a new formal system ZFC* which is, roughly, ZFC augmented by the assumption that ZFC has a countable transitive model M . (X is transitive if for every $x \in X$, every element of x is also an element of X .) The bulk of our work will take place within ZFC*. Specifically, we carry out all forcing arguments in ZFC*; this is where we convert M into a new countable transitive model $M[G]$ in which the desired statement ϕ is also true. Having done this, we then step outside of ZFC* and argue that if ZFC + ϕ were not consistent then ZFC* would not be consistent, and hence ZFC would not be consistent. Thus, consistency of ZFC implies consistency of ZFC + ϕ . This last step requires no set theory and can be formalized in PA.~~

In short, a typical forcing argument is carried out in ZFC* and proves, in ZFC*, that there is a model $M[G]$ of ZFC + ϕ , for some statement ϕ . Given that this result has a proof in ZFC*, we can, using only finitistic methods, draw the conclusion that if ZFC is consistent then so is ZFC + ϕ . The end result is a theorem of Peano arithmetic, but the main part of the proof is a set-theoretic argument in ZFC*.

References

- T. Jech, *Set Theory*, 1978.
K. Kunen, *Set Theory: An Introduction to Independence Proofs*, 1980.
E. Mendelson, *Introduction to Mathematical Logic* (second edition), 1979.

2. Well-ordered sets

This section takes place entirely within ZFC.

Definition 2.1. A well-ordered set is a totally ordered set W with the following property: for any subset $S \subseteq W$, if there exists $a \in W$ such that $x < a$ for all $x \in S$, then there exists a least such a .

In other words, if S has any strict upper bounds then it has a minimal strict upper bound. Or: every subset has an immediate successor, if it has any successors. The significance of this condition is that we can carry out inductive proofs and recursive constructions along W . By the well-ordering property, however far we have carried out a construction there will always be a next step.

Any nonempty well-ordered set W has a smallest element since $\emptyset \subseteq W$ has a least upper bound. If W has more than one element, then the first element has an immediate successor, and so on. If W contains a sequence of elements a_1, a_2, a_3, \dots where a_1 is the smallest element, a_2 is the immediate successor of a_1 , a_3 is the immediate successor of a_2 , etc., then either $W = \{a_n : n \in \mathbb{N}\}$ (and W is order-isomorphic to \mathbb{N}), or else the sequence (a_n) has an immediate successor a_ω . If this does not exhaust W then a_ω must have an immediate successor $a_{\omega+1}$, and so on.

Theorem 2.2. *Let W be a totally ordered set. The following are equivalent:*

- (a) W is well-ordered;
- (b) there is no strictly decreasing sequence in W ;
- (c) every nonempty subset of W has a smallest element.

Proof. (a) \Rightarrow (b). Suppose W is well-ordered and $(a_n) \subseteq W$ is strictly decreasing, i.e., $a_1 > a_2 > \dots$. Define

$$S = \{x \in W : x < a_n \text{ for all } n\}.$$

By the well-ordering property this set has an immediate successor a . Since every a_n is a strict upper bound for S , it follows that $a_n \geq a$ for all n , and since the a_n are decreasing we must then have $a_n > a$ for all n . But then $a \in S$, a contradiction.

(b) \Rightarrow (c). Exercise.

(c) \Rightarrow (a). Suppose (c) holds and let $S \subseteq W$ be a subset which has a strict upper bound. Let T be the set of all strict upper bounds,

$$T = \{y \in W : x < y \text{ for all } x \in S\}.$$

Then (c) implies that T has a smallest element, i.e., S has an immediate successor. This shows that W is well-ordered. ■

Corollary 2.3. *Any subset of a well-ordered set is well-ordered (with the inherited order).*

Proof. Immediate from Theorem 2.2 (b), say. ■

The basic result on well-ordered sets is Theorem 2.7. It requires a definition and two lemmas. In the following “isomorphism” means “order-isomorphism”.

Definition 2.4. Let W be a well-ordered set. An initial segment of W is a set of the form $x^< = \{y \in W : y < x\}$, for some $x \in W$.

Lemma 2.5. *No well-ordered set is isomorphic to an initial segment of itself.*

Proof. Let W be a well-ordered set and let $a \in W$. Suppose $f : W \rightarrow a^<$ is an isomorphism. Then $f(a) < a$, and inductively $f^{n+1}(a) < f^n(a)$ for any $n \in \mathbb{N}$ since f is an isomorphism. Thus the sequence $a > f(a) > f^2(a) > \dots$ is strictly decreasing, which contradicts Theorem 2.2 (b). ■

Lemma 2.6. *Suppose V and W are isomorphic well-ordered sets. Then the isomorphism between them is unique.*

Proof. Given two isomorphisms $f, g : V \rightarrow W$, let $h = g^{-1} \circ f : V \rightarrow V$. For any $x \in V$, the map h establishes an isomorphism between $x^<$ and $h(x)^<$, so Lemma 2.5 implies that $x = h(x)$. This shows that h must be the identity map, and hence $f = g$. ■

Theorem 2.7. *Let V and W be well-ordered sets. Then either (1) V is isomorphic to an initial segment of W ; (2) W is isomorphic to an initial segment of V ; or (3) V and W are isomorphic.*

Proof. Let

$$f = \{\langle x, y \rangle \in V \times W : x^< \text{ is isomorphic to } y^<\}.$$

By Lemma 2.5, for each $x \in V$ there is at most one such $y \in W$, and conversely; thus f is a bijection between a subset of V and a subset of W . It is easy to check that f preserves order.

If $x^<$ is isomorphic to $y^<$ then for any $x_0 < x$ there exists $y_0 < y$ such that $x_0^<$ and $y_0^<$ are isomorphic. Thus, the domain of f is either V or an initial segment of V , and similarly the image of f is either W or an initial segment of W .

If either the domain of f is V or the range of f is W then we are done. So suppose both are initial segments, say $x^< \subseteq V$ and $y^< \subseteq W$. But then f is an isomorphism between $x^<$ and $y^<$, which means the pair $\langle x, y \rangle$ belongs to f , which is a contradiction. ■

Now we prove that well-orderings exist in abundance.

Definition 2.8. Let X be a set all of whose elements are nonempty sets. A choice function for X is a function f with domain X such that $f(x) \in x$ for all $x \in X$.

(Recall that the axiom of choice asserts that every set of nonempty sets has a choice function.)

Theorem 2.9. *Every set can be well-ordered.*

Proof. Let X be a set. By the axiom of choice, let f be a choice function for the set of all nonempty subsets of X . Call a subset $A \subseteq X$ a Zermelo set if there is a well-ordering on A which has the property that every $x \in A$ satisfies $x = f(X - x^<)$. (For instance, \emptyset , $\{f(X)\}$, and $\{f(X), f(X - \{f(X)\})\}$ are all Zermelo sets.)

We first claim that if A and B are isomorphic Zermelo sets, then $A = B$ and the isomorphism is the identity map. To see this let $g : A \rightarrow B$ be an isomorphism, assume g is not the identity map, and let x be the smallest element of A such that $x \neq g(x)$. Then $x^< = g(x)^<$, so

$$x = f(X - x^<) = f(X - g(x)^<) = g(x)$$

(using the fact that A and B are both Zermelo). This is a contradiction, so we conclude that $A = B$ and the isomorphism is the identity map.

Since any initial segment of a Zermelo set is itself a Zermelo set, it now follows from Theorem 2.7 that if A and B are any Zermelo sets, then either A is an initial segment of B , B is an initial segment of A , or $A = B$.

We now claim that the union Z of all Zermelo sets is itself a Zermelo set. Indeed, if x is an element of any Zermelo set A , then by the previous paragraph the initial segment $x^<$ is the same in any Zermelo set that contains x . Thus

$$\{y \in Z : y < x\} = \{y \in A : y < x\}.$$

So the fact that $x = f(X - x^<)$ in A implies the same in Z , and hence Z is Zermelo.

Finally, Z must equal X . Otherwise the set $X - Z$ would be nonempty, and the set $Z \cup \{f(X - Z)\}$ would be a Zermelo set not contained in Z , a contradiction. We conclude that X is well-ordered. ■

Exercises

- (a) Prove Theorem 2.2 (b) \Rightarrow (c).
(b) Let $W_1 \subseteq W_2 \subseteq \dots$ be a nested sequence of well-ordered sets. Assume that for all n the order on W_n agrees with the order it inherits from W_{n+1} . Thus, we obtain a total order on $W = \bigcup_1^\infty W_n$. Is it necessarily a well-ordering?
(c) Work in ZF (all the axioms of ZFC except the axiom of choice). Assuming that every set can be well-ordered, prove the axiom of choice.

3. Ordinals and cardinals

This section takes place entirely within ZFC.

Ordinals are a certain kind of canonical well-ordered set. Informally, every well-ordered set W can be converted into an ordinal in the following way. If W is nonempty then it has a first element; replace it with \emptyset . If W has a second element, replace it with $\{\emptyset\}$. If W has a third element, replace it with $\{\emptyset, \{\emptyset\}\}$. In general, having replaced every element which precedes $x \in W$, replace x with $\{y \in W : y < x\}$. This makes sense both at successor points and at limit points, so we can continue until W is exhausted. We now formalize this assertion.

Definition 3.1. A set X is transitive if for any $x \in X$, every element of x is also an element of X . (Equivalently: $x \in X$ implies $x \subseteq X$.) An ordinal is a transitive set which is well-ordered by \in . We order the ordinals by letting $\alpha < \beta$ if $\alpha \in \beta$. A successor ordinal is an ordinal that has an immediate predecessor, and a limit ordinal is a nonzero ordinal that is not a successor.

Proposition 3.2. *Every well-ordered set is order-isomorphic to an ordinal. If α and β are ordinals then exactly one of $\alpha < \beta$, $\alpha = \beta$, $\alpha > \beta$ is true.*

Proof. Exercise. ■

Every ordinal is the set of all smaller ordinals. Thus, the first few ordinals are

$$\emptyset, \{\emptyset\}, \{\emptyset, \{\emptyset\}\}, \{\emptyset, \{\emptyset, \{\emptyset, \{\emptyset\}\}\}, \dots$$

We set-theoretically encode the natural numbers as finite ordinals by letting $0 = \emptyset$, $1 = \{\emptyset\}$, etc., in general letting $n + 1 = \{0, \dots, n\}$. The first infinite ordinal is

$$\omega = \mathbb{N} = \{0, 1, 2, \dots\}$$

and its immediate successor is $\omega + 1 = \omega \cup \{\omega\}$. (In general the successor of α is $\alpha \cup \{\alpha\}$.) The first uncountable ordinal is denoted ω_1 .

We are working in ZFC, and we now state a “theorem scheme” addressing induction and recursion on the ordinals. For any formulas ϕ and ψ in the language of set theory, the following is a theorem of ZFC.

Theorem 3.3.

(a) *Suppose that for any ordinal α*

$$(\forall \beta < \alpha) \phi(\beta) \rightarrow \phi(\alpha).$$

Then $\phi(\alpha)$ holds for all ordinals α .

(b) *Let a be a set. Suppose that for any nonzero ordinal α and any set x there is a unique set y such that $\psi(\alpha, x, y)$ is true. Then for any ordinal α there is a unique set a_α such that (1) $a_0 = a$ and (2) for all $\alpha > 0$, a_α satisfies $\psi(\alpha, \bigcup_{\beta < \alpha} a_\beta, a_\alpha)$.*

Proof. (a) Suppose $\phi(\alpha)$ fails for some α . Let $S = \{\beta < \alpha : \phi(\beta) \text{ is false}\}$. Then S has a least element β , and $\phi(\gamma)$ holds for all $\gamma < \beta$ but $\phi(\beta)$ is false, contradicting the hypothesis. So $\phi(\alpha)$ must be true for all α .

(b) Uniqueness of the sequence follows from part (a) (induction on α). For existence, let $\phi(\alpha, a, x)$ be the formula “ α is an ordinal and x is a sequence of sets $(a_\beta)_{\beta < \alpha}$ such that $a_0 = a$ and $\psi(\beta, \bigcup_{\gamma < \beta} a_\gamma, a_\beta)$ holds for all $\beta < \alpha$ ”. Then apply part (a) to the formula $(\exists x)\phi(\alpha, a, x)$. ■

For instance, the power set axiom states that every set has a power set, and by the extensionality axiom power sets are unique. Taking $\psi(\alpha, x, y)$ to be “if α is a successor then $y = \mathcal{P}(x)$, and if α is a limit then $y = x$ ” and letting $a = \emptyset$ yields the cumulative hierarchy (V_α) where $V_0 = \emptyset$, $V_{\alpha+1} = \mathcal{P}(V_\alpha)$ for all α , and $V_\alpha = \bigcup_{\beta < \alpha} V_\beta$ when α is a limit ordinal. The rank of a set x is the smallest value of α such that $x \in V_\alpha$. (Exercise: every set belongs to some V_α .)

Definition 3.4. A cardinal is an ordinal that cannot be put in bijection with any smaller ordinal.

Thus every finite ordinal (i.e., $0, 1, 2, \dots$) is a cardinal. Also ω is a cardinal, but $\omega + 1$ is not. Indeed, the next cardinal after ω is ω_1 . We also write $\aleph_0 = \omega$, $\aleph_1 = \omega_1$, and in general $\aleph_\alpha =$ the α th infinite cardinal.

Proposition 3.5. *Every set can be put in bijection with a unique cardinal.*

Proof. Exercise. ■

Definition 3.6. For any set X , $\text{card}(X)$ is the unique cardinal in bijection with X . If λ and κ are cardinals then we define

- (a) $\lambda + \kappa = \text{card}(\lambda \amalg \kappa)$ (the cardinality of their disjoint union),
- (b) $\lambda \cdot \kappa = \text{card}(\lambda \times \kappa)$ (the cardinality of their cartesian product), and
- (c) $\lambda^\kappa = \text{card}(\{\text{all functions from } \kappa \text{ into } \lambda\})$.

To see the reasonableness of the last definition, observe that if m and n are natural numbers then the number of functions from an n -element set into an m -element set is m^n . (Each element of the domain has m possible target values.)

We collect some elementary properties of cardinal arithmetic, with proofs left to the reader.

Proposition 3.7. *Let λ , κ , and θ be cardinals. Then*

- (a) $\lambda + \kappa = \kappa + \lambda$;
- (b) $\lambda + (\kappa + \theta) = (\lambda + \kappa) + \theta$;
- (c) $\lambda \cdot \kappa = \kappa \cdot \lambda$;
- (d) $\lambda \cdot (\kappa \cdot \theta) = (\lambda \cdot \kappa) \cdot \theta$;
- (e) $\lambda \cdot (\kappa + \theta) = \lambda \cdot \kappa + \lambda \cdot \theta$;
- (f) $\lambda^{\kappa+\theta} = \lambda^\kappa \cdot \lambda^\theta$;
- (g) $(\lambda \cdot \kappa)^\theta = \lambda^\theta \cdot \kappa^\theta$.

Observe that every subset of a set X determines a function from X into $2 = \{0, 1\}$, namely its characteristic function; and conversely, every function from X into 2 is the characteristic function of some subset of X .

Thus

$$\begin{aligned} 2^{\text{card}(X)} &= \text{card}(\{\text{functions from } \text{card}(X) \text{ into } 2\}) \\ &= \text{card}(\{\text{functions from } X \text{ into } 2\}) \\ &= \text{card}(\mathcal{P}(X)). \end{aligned}$$

The two basic results on cardinal arithmetic are the following.

Theorem 3.8. *Let κ be a cardinal. Then $\kappa < 2^\kappa$.*

Proof. By the comment preceding the theorem, we must show that for any set X there is an injective map from X into $\mathcal{P}(X)$, but no bijection between X and $\mathcal{P}(X)$. The first task is easy: take $x \in X$ to $\{x\} \in \mathcal{P}(X)$. Now suppose $f : X \rightarrow \mathcal{P}(X)$ is any map. Let $A = \{x \in X : x \notin f(x)\}$. Then for any $x \in X$ we have $x \in f(x)$ if and only if $x \notin A$, so $f(x)$ cannot equal A for any $x \in X$. Thus, there is no map from X onto $\mathcal{P}(X)$. ■

Theorem 3.9. *let κ be an infinite cardinal. Then $\kappa = \kappa^2$.*

Proof. Observe that $\kappa^2 = \kappa \cdot \kappa$. We prove the theorem by induction. We already know that $\aleph_0 = \aleph_0^2$. Now let α be a nonzero ordinal and assume that $\aleph_\beta = \aleph_\beta^2$ for all $\beta < \alpha$. It is easy to see that $\aleph_\alpha \leq \aleph_\alpha^2$. For the reverse inequality, define

$$W = \{\langle x, y \rangle \in \aleph_\alpha^2 : y \leq x\};$$

with the lexicographic order (that is, $\langle x_1, y_1 \rangle < \langle x_2, y_2 \rangle$ if either $x_1 < x_2$ or $x_1 = x_2$ and $y_1 < y_2$) this is a well-ordered set. Any initial segment of W is isomorphic to a subset of γ^2 for some ordinal $\gamma < \aleph_\alpha$, and hence by hypothesis has cardinality strictly less than \aleph_α . Thus, the ordinal that is isomorphic to W (Proposition 3.2) is contained in \aleph_α , and hence $\text{card}(W) \leq \aleph_\alpha$. It easily follows that $\aleph_\alpha^2 \leq \aleph_\alpha$. ■

Corollary 3.10. *For any infinite cardinals λ and κ , we have $\lambda \cdot \kappa = \lambda + \kappa = \max(\lambda, \kappa)$.*

Proof. Without loss of generality suppose $\lambda \geq \kappa$. Then $\lambda \leq \lambda + \kappa \leq 2\lambda \leq \lambda \cdot \kappa \leq \lambda^2 = \lambda$. ■

Corollary 3.11. *For any infinite cardinal λ , we have $2^\lambda = \lambda^\lambda = (2^\lambda)^\lambda$.*

Proof. We have $2^\lambda \leq \lambda^\lambda \leq (2^\lambda)^\lambda = 2^{\lambda^2} = 2^\lambda$. ■

Since $\text{card}(\mathbb{R}) = \text{card}(\mathcal{P}(\mathbb{N})) = 2^{\aleph_0}$, Theorem 3.8 implies that $2^{\aleph_0} \geq \aleph_1$. The continuum hypothesis (CH) is the assertion that $2^{\aleph_0} = \aleph_1$. We will see that CH cannot be decided in ZFC.

Exercises

- Prove Proposition 3.2.
- Let α be an ordinal. Prove that the successor ordinal is $\alpha \cup \{\alpha\}$.
- Prove that every set belongs to some V_α .
- Prove Proposition 3.5.
- Use cardinal arithmetic to prove that the cardinality of the set of all countable subsets of \mathbb{R} equals the cardinality of \mathbb{R} .

4. Relativization, reflection, and collapse

This section takes place entirely within ZFC.

We would like to have sets that can play the role of “miniature universes”. In principle any set M can be thought of this way, but the better closure properties M has the more realistically it will mimic the real universe of sets. Ideally we would like all the axioms of ZFC to be “true in M ”. What this means is explained in the following definition.

Definition 4.1. Let M be a set and let ϕ be a formula in the language of set theory. The relativization of ϕ to M , denoted ϕ^M , is the formula obtained from ϕ by replacing every quantifier with its restriction to M . That is, replace every $(\forall x)$ with $(\forall x \in M)$ and replace every $(\exists x)$ with $(\exists x \in M)$. If a set S is defined by proving

$$(\exists! x)\phi(x)$$

where ϕ has only one free (unquantified) variable x , and letting S be the unique x satisfying $\phi(x)$, then we can define the relativization of S to M by proving

$$(\exists! x \in M)\phi^M(x)$$

and letting S^M be the unique $x \in M$ satisfying $\phi^M(x)$. If there are no free variables in ϕ (it is a sentence), then we say that ϕ is true in M or that M satisfies ϕ , written $M \models \phi$, if ϕ^M is true. More generally, if the free variables of ϕ are among x_1, \dots, x_n and $u_1, \dots, u_n \in M$ then we say that $\phi(u_1, \dots, u_n)$ is true in M and write $M \models \phi(u_1, \dots, u_n)$ if $\phi^M(u_1, \dots, u_n)$ is true.

Note that the definition of S^M is ambiguous, since the set S might be definable using another formula ψ whose relativization to M is not equivalent to ϕ^M . Also note that in general S^M need not exist. But both of these problems effectively disappear if all the axioms of ZFC are true in M , since then (1) two formulas that are provably equivalent outside M are also provably equivalent inside M , and (2) if we can prove the existence of S outside M then we can prove the existence of S^M inside M .

Next we take up the problem of finding sets M in which various formulas ϕ are true. We do this by the reflection principle (Theorem 4.3).

For any finite list of formulas ϕ_1, \dots, ϕ_n whose free variables are among x, x_1, \dots, x_k , the following is provable in ZFC.

Lemma 4.2. *Let M_0 be a countable set. Then there is a countable set M which contains M_0 , such that for every $u_1, \dots, u_k \in M$ we have*

$$\begin{array}{ccc} (\exists x)\phi_1(x, u_1, \dots, u_k) & \text{implies} & (\exists x \in M)\phi_1(x, u_1, \dots, u_k). \\ & & \vdots \\ (\exists x)\phi_n(x, u_1, \dots, u_k) & \text{implies} & (\exists x \in M)\phi_n(x, u_1, \dots, u_k). \end{array}$$

Proof. The set of k -tuples $\langle u_1, \dots, u_k \rangle \in M_0^k$ is countable. Construct a set M_1 by adding to M_0 , for each i and each such k -tuple, a value of x such that $\phi_i(x, u_1, \dots, u_k)$, if any such value exists. Then M_1 is also countable, and hence so is the set of k -tuples in M_1 . Construct a set M_2 by adding to M_1 , for each i and each such k -tuple, a value of x such that $\phi_i(x, u_1, \dots, u_k)$, if any such value exists. Continue recursively and let $M = \bigcup_0^\infty M_j$. Then M is a countable union of countable sets and hence is countable. Now suppose that for some $u_1, \dots, u_k \in M$ there exists x such that $\phi_i(x, u_1, \dots, u_k)$. Then u_1, \dots, u_k lie in M_j for some j , and by the construction of M_{j+1} there exists $x \in M_{j+1}$ satisfying $\phi_i(x, u_1, \dots, u_k)$. So there exists such an x in M . ■

For any finite list of sentences ϕ_1, \dots, ϕ_n which are provable in ZFC, the following is a theorem of ZFC.

Theorem 4.3. *Let M_0 be a countable set. Then there is a countable set M which contains M_0 , such that $M \models \phi_1 \wedge \dots \wedge \phi_n$.*

Proof. Let ψ_1, \dots, ψ_m be a list of formulas that contains the sentences ϕ_1, \dots, ϕ_n and is subformula closed, meaning that if $\neg\phi$ is in the list, so is ϕ ; if $\phi \wedge \psi$ is in the list, so are ϕ and ψ ; if $(\forall x)\phi$ is in the list, so is ϕ , etc. We can no longer assume the ψ_i are sentences. We will find M such that that for each ψ_i in the list with free variables x_1, \dots, x_k , we have

$$\psi_i(u_1, \dots, u_k) \leftrightarrow \psi_i^M(u_1, \dots, u_k) \tag{*}$$

for any $u_1, \dots, u_k \in M$. If ψ_i has no free variables, this means $\psi_i \leftrightarrow \psi_i^M$, so the truth of ϕ_1, \dots, ϕ_n implies the truth of $\phi_1^M, \dots, \phi_n^M$.

Apply Lemma 4.2 to ψ_1, \dots, ψ_m to obtain the desired set M . We prove (*) by induction on the complexity of the formulas ψ_i . If ψ_i is atomic then (*) is trivial. If (*) holds for ϕ and ψ then it trivially holds for $\neg\phi$, $\phi \wedge \psi$, and $\phi \vee \psi$. Finally, the content of Lemma 4.2 is that if (*) holds for ϕ then it holds for $(\exists x)\phi$. This completes the proof. (We need not consider universal quantification separately, since $(\forall x)\phi$ is logically equivalent to $\neg(\exists x)\neg\phi$.) ■

In particular, if ϕ_1, \dots, ϕ_n are among the axioms of ZFC then we can prove, in ZFC, the existence of a countable set M in which ϕ_1, \dots, ϕ_n are true. We say that there is a model of any finite fragment of ZFC. This result is best possible since, assuming ZFC is consistent, we cannot prove in ZFC that ZFC has a model — this is a consequence of Gödel’s second incompleteness theorem. But merely having a model of any finite fragment is a very strong result, since any theorem of ZFC is proven using only finitely many axioms. Hence any theorem of ZFC, or any finitely many theorems of ZFC, are provably true in some countable set M .

It is convenient to work with models that are transitive (see Definition 3.1), and this can always be arranged provided M is extensional, meaning that for any distinct $x, y \in M$ there exists $u \in M$ either in x but not y or in y but not x . (Equivalently: the axiom of extensionality is true when relativized to M .) The technique that achieves transitivity is called Mostowski collapse. This is essentially a generalization of the fact that every well-ordered set is order-isomorphic to an ordinal (Proposition 3.2).

Theorem 4.4. *Let M be an extensional set. Then there is a transitive set N and a bijection $f : M \rightarrow N$ such that $x \in y$ in M if and only if $f(x) \in f(y)$ in N .*

Proof. Say that a set $M_0 \subseteq M$ is Mostowski if it is downward closed (i.e., if $x \in M_0$ then any $u \in M$ that is in x is in M_0) and there is a transitive set N_0 and a bijection $f_0 : M_0 \rightarrow N_0$ such that $x \in y$ in M_0 if and only if $f_0(x) \in f_0(y)$ in N_0 .

Exercise: if M_0 is Mostowski then the N_0 and f_0 which verify this fact are unique. The intersection of any two Mostowski sets is Mostowski, and the corresponding bijections with transitive sets agree on the intersection. Any union of Mostowski sets is Mostowski.

Let M' be the union of all Mostowski sets. It follows from the preceding paragraph that M' is Mostowski. Let $f : M' \rightarrow N$ be the bijection that verifies this and suppose $M' \neq M$. By the axiom of foundation there is an \in -minimal element x of $M - M'$. Let $y = \{f(u) : u \in x \cap M\}$. It is then straightforward to verify that $M' \cup \{x\}$ is Mostowski via a bijection with $N \cup \{y\}$, contradicting maximality of M' . Thus $M' = M$, which verifies the theorem. ■

The fact that M and N are isomorphic in the sense of the theorem means that ϕ^M will be true if and only if ϕ^N is true, for any sentence ϕ . Thus, combining the reflection principle with Mostowski collapse yields the next corollary. For any finite list of sentences ϕ_1, \dots, ϕ_n which are provable in ZFC, the following is provable in ZFC.

Corollary 4.5. *There is a countable transitive set M such that $M \models \phi_1 \wedge \dots \wedge \phi_n$.*

Exercises

- (a) In the proof of Theorem 4.3, make explicit the argument that if $(*)$ holds for ϕ then it holds for $(\exists x)\phi$.
- (b) In the proof of Theorem 4.4, prove that if M_0 is Mostowski then the N_0 and f_0 which verify this fact are unique. (Hint: use the axiom of foundation.) Prove that the intersection of any two Mostowski sets is Mostowski, and the corresponding bijections with transitive sets agree on the intersection. Use this to show that any union of Mostowski sets is Mostowski.

5. Finitistic consistency proofs

This section is finitistically valid. All results could be formalized in PA.

All forcing arguments will take place in the formal system ZFC*, which we define now.

Definition 5.1. The language of ZFC* is the language of ZFC together with one constant symbol M . The atomic formulas of ZFC* are all formulas of the form $x \in y$ or $x = y$, where x and y can be replaced by any variables, plus all such formulas in which either or both variables are replaced by M . The formulas of ZFC* are built up from the atomic formulas in the usual way. The axioms of ZFC* come in four categories:

- (a) every axiom of ZFC is also an axiom of ZFC*;
- (b) $(\forall x)\phi(x) \rightarrow \phi(M)$ is an axiom of ZFC* for every formula ϕ and any variable in place of x ;
- (c) the statement “M is countable and transitive” is an axiom of ZFC*;
- (d) the relativization of any axiom of ZFC to M is an axiom of ZFC*.

In effect, ZFC* is ZFC augmented by the fact that there is a countable transitive set M which models ZFC. However, the way this fact is formalized is slightly subtle. Working in ZFC*, we have a countable transitive set M, and for any axiom ϕ of ZFC we know that ϕ is true in M. But (assuming ZFC* is consistent) we cannot prove in ZFC* a single statement to the effect that all axioms of ZFC are true in M. This is good because if we could do this, then in ZFC* we could prove the consistency of ZFC, and Gödel's second incompleteness theorem would then prevent us from being able to prove the following result.

Lemma 5.2. *If ZFC is consistent, so is ZFC*.*

Proof. Suppose ZFC* is inconsistent. Then there is a proof of $\phi \wedge \neg\phi$ for some formula ϕ . This proof involves only finitely many axioms of type (d) in Definition 5.1, say $\phi_1^M, \dots, \phi_n^M$ where ϕ_1, \dots, ϕ_n are axioms of ZFC.

By Corollary 4.5, we can prove in ZFC that there is a countable transitive set M such that $\phi_1^M, \dots, \phi_n^M$ are true. Having done this, we can then copy in ZFC the original proof of $\phi \wedge \neg\phi$ in ZFC*, everywhere replacing M with M . The result is a proof in ZFC of a contradiction. Thus, we have shown that any inconsistency in ZFC* could be mechanically converted into an inconsistency in ZFC. ■

In a typical forcing argument, we will work in ZFC* and enlarge the set M to another countable transitive set $M[G]$, in such a way that we ensure that some special formula ϕ (e.g., the continuum hypothesis) is true in $M[G]$. For any axiom ψ of ZFC, we will also be able to show that ψ is true in $M[G]$. (But again, we cannot prove a single statement to the effect that all axioms of ZFC are true in $M[G]$.) This gives us the hypothesis of the following theorem.

Theorem 5.3. *Let ϕ be a formula in the language of ZFC. Suppose that in ZFC* we can define a countable transitive set N such that (1) in ZFC* we can prove ϕ^N and (2) for any axiom ψ of ZFC, in ZFC* we can prove ψ^N . Then if ZFC is consistent, so is ZFC + ϕ .*

Proof. Suppose ZFC + ϕ is not consistent. Then there is a proof in ZFC + ϕ of some contradiction. Since in ZFC* we can prove the relativization of ϕ to N as well as the relativization of any axiom of ZFC to N , we can relativize the entire proof of the contradiction to N and this becomes a valid proof in ZFC* of a contradiction. So ZFC* is not consistent. By the lemma, it then follows that ZFC is not consistent. Thus, we have shown that any inconsistency in ZFC + ϕ could be mechanically converted into an inconsistency in ZFC. ■

6. Generic extensions

From this point on, except in applications sections we work in ZFC.*

We have a countable transitive set M in which the axioms of ZFC are true. We would like to enlarge it by adding in a structure that makes some desired additional assertion true. For example, we may want to add a bijection between $\mathcal{P}(\mathbb{N})$ and \aleph_1 in order to make the continuum hypothesis true, or a set of \aleph_α real numbers for some $\alpha > 1$ to make the continuum hypothesis false. In order for this procedure to work it will be important that the structure being added has no “special properties” relative to the model M — it is generic. We ensure this by defining, in M, a partially ordered set of all possible partial constructions of the desired structure. As we move down this partially ordered set we specify the desired structure more precisely. We can then identify a “generic” path down the poset which describes a “generic” structure of the desired type.

Definition 6.1. A notion of forcing is a preordered set $P \in M$ (i.e., P is equipped with a relation \leq which is reflexive and transitive) with greatest element 1_P .

- (a) If $p, q \in P$ and $p \leq q$ then p is an extension of q . Two elements are compatible if they have a common extension.
- (b) A subset D of P is dense if every $p \in P$ has an extension in D . It is dense below p if every $q \leq p$ has an extension in D .
- (c) A filter of P is a subset G which is upwards closed (if $p \in G$ and $p \leq q$ then $q \in G$) and directed downwards (every pair of elements of G have a common extension in G). It is generic if $G \cap D \neq \emptyset$ for every dense subset $D \in \mathcal{M}$.

Note that in the definition of generic we do not require G to meet every dense subset, only those that appear in \mathcal{M} . Indeed, suppose P satisfies the mild condition that every element of P lies above a pair of incompatible elements; then the complement of any filter is dense (exercise), so no filter meets all dense subsets. In fact, this shows that under this condition on P , no generic filter can lie in \mathcal{M} (since if it did its complement would also lie in \mathcal{M} and this would prevent it from being generic). However, generic filters do always exist.

Lemma 6.2. *Let P be a notion of forcing and let $p_0 \in P$. Then there is a generic filter of P that contains p_0 .*

Proof. Since \mathcal{M} is countable, we can enumerate all dense subsets $D \subseteq P$ which lie in \mathcal{M} . Let (D_n) be such an enumeration. Recursively define a sequence (p_n) by choosing $p_{n+1} \leq p_n$ in D_n ; we can always find such an element p_{n+1} by density of D_n . Finally, let

$$G = \{p \in P : p_n \leq p \text{ for some } n\}.$$

It is immediate that G is upwards closed and that it meets every dense subset in \mathcal{M} . To see that it is directed downward, let $p, q \in G$. Then $p_m \leq p$ and $p_n \leq q$ for some m, n . Then whichever of p_m and p_n is smaller is a common extension of p and q in G . Thus G is a generic filter. ■

For the remainder of this section, fix a notion of forcing P and a generic filter G . Our next goal is to define the set $\mathcal{M}[G]$. This is supposed to be the smallest transitive model of ZFC which is larger than \mathcal{M} and contains the generic filter G . We construct $\mathcal{M}[G]$ by first specifying names of all of its elements. These names actually lie in \mathcal{M} , but we need to use G to determine the sets that the names identify.

Definition 6.3. P -names are defined recursively by rank. Let $\mathcal{P}_0 = \{\emptyset\}$, and for any ordinal $\alpha > 0$ in \mathcal{M} let \mathcal{P}_α consist of all $\tau \in \mathcal{M}$ such that τ is a set of ordered pairs of the form $\langle \sigma, p \rangle$ with $p \in P$ and $\sigma \in \mathcal{P}_\beta$ for some $\beta < \alpha$. A P -name is any element of any \mathcal{P}_α , and the name rank of a P -name τ is the least α such that $\tau \in \mathcal{P}_\alpha$. The domain of a P -name τ , $\text{dom}(\tau)$, is the set of σ such that $\langle \sigma, p \rangle \in \tau$ for some $p \in P$.

Note that since $P \in \mathcal{M}$ and \mathcal{M} satisfies the axioms of ZFC, the preceding definition makes sense in \mathcal{M} , i.e., for each α we have $\mathcal{P}_\alpha \in \mathcal{M}$. (Use Theorem 3.3 (b).)

Definition 6.4. The value of a P -name τ is defined recursively on name rank by

$$\text{val}_G(\tau) = \{\text{val}_G(\sigma) : \langle \sigma, p \rangle \in \tau \text{ for some } p \in G\}$$

and $\mathcal{M}[G]$ is the set of values of all P -names.

When there is only one generic filter in play, we will usually write $\text{val}(\tau)$ instead of $\text{val}_G(\tau)$.

Definition 6.5. The P -name \check{x} is defined recursively on rank for all $x \in \mathcal{M}$ by

$$\check{x} = \{\langle \check{u}, 1_P \rangle : u \in x\}.$$

The P -name Γ is defined by

$$\Gamma = \{\langle \check{p}, p \rangle : p \in P\}.$$

Proposition 6.6. $\mathcal{M} \subseteq \mathcal{M}[G]$ and $G \in \mathcal{M}[G]$. $\mathcal{M}[G]$ is countable and transitive.

Proof. The first statement holds because $\text{val}(\dot{x}) = x$ for all $x \in M$ and $\text{val}(\Gamma) = G$ (exercise). $M[G]$ is countable because the set of P -names is a subset of M and hence is countable. $M[G]$ is transitive because the value of any P -name (a typical element of $M[G]$) is by definition a set of values of P -names (which are themselves elements of $M[G]$). ■

We want to check that the axioms of ZFC are true in $M[G]$. Some of them are easy, and do not even use the fact that G is generic:

Proposition 6.7. *The axioms of extensionality, pairing, infinity, and foundation are true in $M[G]$.*

Proof. Extensionality follows from the fact that $M[G]$ is transitive. Pairing holds because if σ and τ are any two P -names then $\{\langle \sigma, 1_P \rangle, \langle \tau, 1_P \rangle\}$ is a P -name whose value is $\{\text{val}(\sigma), \text{val}(\tau)\}$. Infinity follows from the fact that $M \subseteq M[G]$. Foundation is trivial (it is true in any set). ■

Exercises

- (a) Let G be a generic filter, let $p \in G$, and suppose $D \in M$ is dense below p . Show that $G \cap D \neq \emptyset$.
- (b) Let P be a notion of forcing in which every element lies above a pair of incompatible elements. Prove that the complement of any filter is dense.
- (c) Show that $\text{val}(\dot{x}) = x$ for all $x \in M$ and $\text{val}(\Gamma) = G$.
- (d) Prove that every ordinal in $M[G]$ is in M .

7. The fundamental theorem of forcing

Throughout this section, fix a notion of forcing P .

What makes the $M[G]$ construction a workable technique for proving independence results is our ability to control properties of $M[G]$, for any generic filter G , using P . The key concept is the following. (This is a definition scheme, depending on a formula ϕ with free variables among x_1, \dots, x_n .)

Definition 7.1. Let τ_1, \dots, τ_n be P -names and let $p \in P$. Then p forces $\phi(\tau_1, \dots, \tau_n)$ if

$$M[G] \models \phi(\text{val}_G(\tau_1), \dots, \text{val}_G(\tau_n))$$

for all generic filters G that contain p . We write $p \Vdash \phi(\tau_1, \dots, \tau_n)$.

Observe that if $p \Vdash \phi(\tau_1, \dots, \tau_n)$ and $q \leq p$ then $q \Vdash \phi(\tau_1, \dots, \tau_n)$, because any generic filter that contains q also contains p .

We now state the fundamental theorem of forcing. It essentially says that (a) anything that is true in $M[G]$ is forced by some $p \in G$ and (b) the forcing relation can be determined within M . For each formula ϕ with free variables among x_1, \dots, x_n there is a formula ψ such that the following is provable in ZFC*.

Theorem 7.2. *Let P be a notion of forcing.*

- (a) *Let τ_1, \dots, τ_n be P -names. For any generic filter G , if $M[G] \models \phi(\text{val}_G(\tau_1), \dots, \text{val}_G(\tau_n))$ then some $p \in G$ forces $\phi(\tau_1, \dots, \tau_n)$.*
- (b) *Let $p, \tau_1, \dots, \tau_n \in M$. Then $M \models \psi(p, \tau_1, \dots, \tau_n)$ if and only if τ_1, \dots, τ_n are P -names, $p \in P$, and $p \Vdash \phi(\tau_1, \dots, \tau_n)$.*

Proof. We will simultaneously prove (a) and (b) in the case that ϕ is $x_1 \subseteq x_2$. For part (b), let \mathcal{F}_α be the set of all triples $\langle p, \tau_1, \tau_2 \rangle$ such that τ_1 and τ_2 are P -names of name rank at most α and $p \in P$ forces $\tau_1 \subseteq \tau_2$. It will follow from (ii) below that \mathcal{F}_α can be constructed in M from $\bigcup_{\beta < \alpha} \mathcal{F}_\beta$. This implies that the sequence (\mathcal{F}_α) is definable in M (cf. Theorem 3.3 (b)), and ψ can then be taken to be a formalization of the assertion “ $\langle p, \tau_1, \tau_2 \rangle \in \mathcal{F}_\alpha$ for some ordinal α ”.

The proof goes by induction on the name ranks of τ_1 and τ_2 . Assume that for all P -names τ_1 and τ_2 of name rank less than α we know

- (i) for any generic filter G , if $\text{val}_G(\tau_1) \subseteq \text{val}_G(\tau_2)$ then some $p \in G$ forces $\tau_1 \subseteq \tau_2$;
- (ii) $p \in P$ forces $\tau_1 \subseteq \tau_2$ if and only if
 - (†) for every $\langle \pi_1, s_1 \rangle \in \tau_1$, every element of P that is less than both p and s_1 lies above some $q \in P$ such that $q \leq s_2$ and q forces $\pi_1 = \pi_2$, for some $\langle \pi_2, s_2 \rangle \in \tau_2$.

We will verify that these also hold when one or both of τ_1 and τ_2 are P -names of name rank α . The truth of the theorem in the case that ϕ is $x_1 \subseteq x_2$ will then follow by induction on α .

First let τ_1 and τ_2 be P -names of name rank less than or equal to α , let G be a generic filter, and suppose $\text{val}(\tau_1) \subseteq \text{val}(\tau_2)$; we want to show that some $p \in G$ forces $\tau_1 \subseteq \tau_2$. Let D be the set of $r \in P$ such that

- (‡) there exists $\langle \pi_1, s_1 \rangle \in \tau_1$ with $s_1 \geq r$ and such that for any $\langle \pi_2, s_2 \rangle \in \tau_2$ and any $q \leq s_2$ which forces $\pi_1 = \pi_2$, r and q are incompatible.

The induction hypothesis for (ii) implies that in the case that ϕ is $x_1 \subseteq x_2$, the forcing relation for P -names of name rank less than α can be determined in M . Using the fact that $x_1 = x_2$ if and only if $x_1 \subseteq x_2$ and $x_2 \subseteq x_1$, we see that the condition that q forces $\pi_1 = \pi_2$ in (‡) is determined in M , and it follows that $D \in M$. Then since G is generic, either some $r \in D$ belongs to G or else there is an element of G none of whose extensions lie in D . We claim that the first case cannot obtain. Suppose to the contrary that some $r \in G$ satisfies (‡) and let $\langle \pi_1, s_1 \rangle \in \tau_1$ verify this. Since $r \in G$ and $r \leq s_1$, it follows that $\text{val}(\pi_1) \in \text{val}(\tau_1) \subseteq \text{val}(\tau_2)$, so there must exist $\langle \pi_2, s_2 \rangle \in \tau_2$ such that $s_2 \in G$ and $\text{val}(\pi_1) = \text{val}(\pi_2)$. But since π_1 and π_2 are both P -names of name rank less than α , the induction hypothesis for (i) yields that some $p \in G$ forces $\pi_1 = \pi_2$. (Some $p_1 \in G$ forces $\pi_1 \subseteq \pi_2$ and some $p_2 \in G$ forces $\pi_2 \subseteq \pi_1$; find $p \in G$ less than both p_1 and p_2 .) Since G is a filter we can find $q \in G$ which is less than both p and s_2 , and then q forces $\pi_1 = \pi_2$, but q is compatible with r since both lie in G . This contradicts (‡), so the claim is proven.

We now know that there exists $p \in G$ such that every $r \leq p$ fails to satisfy (‡). We will show that p forces $\tau_1 \subseteq \tau_2$. Let G' be any generic filter that contains p and let $\langle \pi_1, s_1 \rangle$ be any element of τ_1 such that $s_1 \in G'$. Since G' is a filter there exists $r \in G'$ such that $r \leq p$ and $r \leq s_1$; by the failure of (‡), the set of $q \in P$ that force $\pi_1 = \pi_2$ for some $\langle \pi_2, s_2 \rangle \in \tau_2$ with $s_2 \geq q$ is dense below r . Since G' is generic it follows that G' must contain some such q , and therefore $\text{val}_{G'}(\pi_1) = \text{val}_{G'}(\pi_2) \in \text{val}_{G'}(\tau_2)$. Since $\text{val}_{G'}(\pi_1)$ was an arbitrary element of $\text{val}_{G'}(\tau_1)$, we have shown that $\text{val}_{G'}(\tau_1) \subseteq \text{val}_{G'}(\tau_2)$. We conclude that p forces $\tau_1 \subseteq \tau_2$, which completes the proof of (i) in the case that ϕ is $x_1 \subseteq x_2$.

To prove (ii), we must show that p forces $\tau_1 \subseteq \tau_2$ if and only if (†) holds. The reverse direction is an exercise. For the forward direction, suppose p forces $\tau_1 \subseteq \tau_2$, fix $\langle \pi_1, s_1 \rangle \in \tau_1$, and suppose $r \leq p$ and $r \leq s_1$. Let G' be any generic filter that contains r . Then $\text{val}_{G'}(\tau_1) \subseteq \text{val}_{G'}(\tau_2)$ since $r \leq p$, so $\text{val}_{G'}(\pi_1)$ must equal $\text{val}_{G'}(\pi_2)$ for some P -name π_2 such that $\langle \pi_2, s_2 \rangle \in \tau_2$ and $s_2 \in G'$. The induction hypothesis for (i) then implies that there exists $q \in G'$ which forces $\pi_1 = \pi_2$. (Some q_1 forces $\pi_1 \subseteq \pi_2$ and some q_2 forces $\pi_2 \subseteq \pi_1$; find $q \in G'$ less than both q_1 and q_2 .) Since $r, s_2 \in G'$, we may take q less than both r and s_2 , which verifies (†). This completes the proof of (ii) in the case that ϕ is $x_1 \subseteq x_2$.

We have proven the theorem in the case that ϕ is $x_1 \subseteq x_2$. This easily implies the result in the case that ϕ is $x_1 = x_2$. The remainder of the proof involves proving versions of (i) and (ii) first when ϕ is the atomic formula $x_1 \in x_2$, and then inductively for formulas ϕ of greater complexity. We omit details; the case when ϕ is $x_1 = x_2$ is by far the most difficult. ■

For any formulas ϕ and ψ with free variables among x_1, \dots, x_n the following is provable in ZFC*.

Corollary 7.3. *Let τ_1, \dots, τ_n be P -names.*

- (a) *The set of $p \in P$ which either force $\phi(\tau_1, \dots, \tau_n)$ or $\neg\phi(\tau_1, \dots, \tau_n)$ is dense.*
- (b) *$p \Vdash \neg\phi(\tau_1, \dots, \tau_n)$ if and only if no $q \leq p$ forces $\phi(\tau_1, \dots, \tau_n)$.*
- (c) *$p \Vdash (\phi \wedge \psi)(\tau_1, \dots, \tau_n)$ if and only if $p \Vdash \phi(\tau_1, \dots, \tau_n)$ and $p \Vdash \psi(\tau_1, \dots, \tau_n)$.*
- (d) *$p \Vdash (\forall x)\phi(x, \tau_2, \dots, \tau_n)$ if and only if $p \Vdash \phi(\sigma, \tau_2, \dots, \tau_n)$ for every P -name σ .*

Proof. For (a), let $q \in P$ and find a generic filter G that contains q . Then either

$$M[G] \models \phi(\text{val}(\tau_1), \dots, \text{val}(\tau_n)) \quad \text{or} \quad M[G] \models \neg\phi(\text{val}(\tau_1), \dots, \text{val}(\tau_n)),$$

so by Theorem 7.2 (a) there exists $r \in G$ which forces either $\phi(\tau_1, \dots, \tau_n)$ or $\neg\phi(\tau_1, \dots, \tau_n)$. Since G is a filter there exists $p \in G$ such that $p \leq q$ and $p \leq r$; thus there is an extension of q that forces one of the two statements, which verifies density. The remaining parts are left as exercises. ■

Exercises

- (a) In the proof of Theorem 7.2, show that (†) implies $p \Vdash \tau_1 \subseteq \tau_2$.
- (b) Prove Corollary 7.3 (b), (c), and (d).

8. $M[G]$ models ZFC

We are now in a position to check that the remaining axioms of ZFC hold in $M[G]$. Recall that we have already verified extensionality, pairing, infinity, and foundation (Proposition 6.7). For any formula ϕ the following is a theorem of ZFC*.

Theorem 8.1. *Let P be a notion of forcing and let G be a generic filter of P . Then the axioms of union, power set, and choice are true in $M[G]$, as are the instances of separation and replacement involving the formula ϕ .*

Proof. Rather than verify all of the axioms we choose two representative cases: separation and power set. For the axiom of separation, let σ and τ_1, \dots, τ_n be P -names; we need to show that

$$S = \{a \in \text{val}(\sigma) : M[G] \models \phi(a, \text{val}(\tau_1), \dots, \text{val}(\tau_n))\}$$

is in $M[G]$. We do this by exhibiting the P -name

$$\rho = \{ \langle \pi, p \rangle \in \text{dom}(\sigma) \times P : p \Vdash \phi(\pi, \tau_1, \dots, \tau_n) \}.$$

This is a legitimate P -name; specifically, it belongs to M because the forcing relation is determined within M (Theorem 7.2 (b)).

We will show that $\text{val}(\rho) = S$. First, let a be an arbitrary element of $\text{val}(\rho)$, so that there exists $\langle \pi, p \rangle \in \rho$ such that $p \in G$ and $a = \text{val}(\pi)$. By the definition of ρ we then have that $a = \text{val}(\pi) \in \text{val}(\sigma)$ and $M[G] \models \phi(\text{val}(\pi), \text{val}(\tau_1), \dots, \text{val}(\tau_n))$, and this immediately yields $a \in S$. For the reverse direction, let $a \in \text{val}(\sigma)$ and suppose $\phi(a, \text{val}(\tau_1), \dots, \text{val}(\tau_n))$ is true in $M[G]$. Then $a = \text{val}(\pi)$ for some $\pi \in \text{dom}(\sigma)$. By Theorem 7.2 (a), there exists $p \in G$ which forces “ $\pi \in \sigma$ and $\phi(\pi, \tau_1, \dots, \tau_n)$ ”, so $\langle \pi, p \rangle \in \rho$ and hence $a = \text{val}(\pi) \in \text{val}(\rho)$. This completes the proof of separation.

For the power set axiom, let σ be any P -name and define

$$\rho = \{ \langle \tau, 1_P \rangle : \tau \in \mathcal{P}(\text{dom}(\sigma) \times P)^M \}$$

(i.e., τ ranges over all subsets of $\text{dom}(\sigma) \times P$ that belong to M). We will show that every subset of $\text{val}(\sigma)$ in $M[G]$ belongs to $\text{val}(\rho)$. This is enough, since $\text{val}(\rho) \in M[G]$ and an appropriate instance of separation will then show $\mathcal{P}(\text{val}(\sigma))^{M[G]} \in M[G]$. To prove this let μ be a P -name such that $\text{val}(\mu) \subseteq \text{val}(\sigma)$. Define

$$\tau = \{ \langle \pi, p \rangle : \pi \in \text{dom}(\sigma) \text{ and } p \Vdash \pi \in \mu \}.$$

It is clear that $\langle \tau, 1_P \rangle \in \rho$, so that $\text{val}(\tau) \in \text{val}(\rho)$. Finally, we check that $\text{val}(\mu) = \text{val}(\tau)$. In one direction, any element of $\text{val}(\mu)$ is of the form $\text{val}(\pi)$ for some $\pi \in \text{dom}(\sigma)$. Since $\text{val}(\pi)$ is in $\text{val}(\mu)$ some $p \in G$ forces $\pi \in \mu$, so that $\langle \pi, p \rangle \in \tau$ and hence $\text{val}(\pi) \in \text{val}(\tau)$. In the reverse direction, any element of $\text{val}(\tau)$ is of the

form $\text{val}(\pi)$ such that $\langle \pi, p \rangle \in \tau$ for some $p \in G$ with $p \Vdash \pi \in \mu$, and this implies that $\text{val}(\pi) \in \text{val}(\mu)$. This completes the proof that $\text{val}(\mu) = \text{val}(\tau)$. \blacksquare

The preceding result motivates the following definition scheme. Let ϕ be a formula with free variables among x_1, \dots, x_n .

Definition 8.2. ϕ is absolute if for any notion of forcing P and any generic filter G of P , we have

$$\mathbb{M} \models \phi(u_1, \dots, u_n) \quad \text{if and only if} \quad \mathbb{M}[G] \models \phi(u_1, \dots, u_n)$$

for any $u_1, \dots, u_n \in \mathbb{M}$. Also, recalling the caveat given after Definition 4.1, we say that (the standard definition of) a set S is absolute if $S^{\mathbb{M}} = S^{\mathbb{M}[G]}$ for any notion of forcing P and any generic filter G of P .

For example, Proposition 6.7 and Theorem 8.1 show that every axiom of ZFC is absolute. Also, the formula “ x is an ordinal” is absolute, as is \aleph_0 . But we will see that “ x is a cardinal” and \aleph_1 are not absolute — for some notions of forcing, the ordinal $\aleph_1^{\mathbb{M}}$ becomes countable in $\mathbb{M}[G]$.

Exercises

(a) Prove that the union axiom holds in $\mathbb{M}[G]$, for any notion of forcing P and any generic filter G of P .

9. Forcing CH

In this section we will consider the following notion of forcing P . The elements of P are all the functions f in \mathbb{M} such that $\mathbb{M} \models$ “ f is a bijection between countable subsets of $\mathcal{P}(\mathbb{N})$ and \aleph_1 ”. That is, f is a bijection in \mathbb{M} between subsets of $\mathcal{P}(\mathbb{N})^{\mathbb{M}}$ and $\aleph_1^{\mathbb{M}}$ which are countable according to \mathbb{M} . (Of course, since \mathbb{M} is countable both $\mathcal{P}(\mathbb{N})^{\mathbb{M}}$ and $\aleph_1^{\mathbb{M}}$ are actually themselves countable. But $\mathbb{M} \models$ “ $\mathcal{P}(\mathbb{N})$ and \aleph_1 are uncountable”.) We call elements of P countable partial bijections between $\mathcal{P}(\mathbb{N})^{\mathbb{M}}$ and $\aleph_1^{\mathbb{M}}$.

We order P by setting $f \leq g$ if $\text{dom}(g) \subseteq \text{dom}(f)$ and $f|_{\text{dom}(g)} = g$. That is, f extends g as an element of P if and only if f extends g as a function.

We will show that the continuum hypothesis is true in $\mathbb{M}[G]$, for any generic filter G of P . The idea is that elements of P give us partial information on how to construct a bijection between $\mathcal{P}(\mathbb{N})^{\mathbb{M}}$ and $\aleph_1^{\mathbb{M}}$. When we pass to $\mathbb{M}[G]$, it will be easy to see that an actual bijection between them has been introduced. The more subtle point is that we must also verify that $\mathcal{P}(\mathbb{N})^{\mathbb{M}[G]} = \mathcal{P}(\mathbb{N})^{\mathbb{M}}$ and $\aleph_1^{\mathbb{M}[G]} = \aleph_1^{\mathbb{M}}$; if passing to $\mathbb{M}[G]$ introduces new sets of natural numbers, for example, then merely having a bijection between the sets that were $\mathcal{P}(\mathbb{N})$ and \aleph_1 in \mathbb{M} would not verify the continuum hypothesis in $\mathbb{M}[G]$. The key lemma is the following.

Definition 9.1. A preordered set P is ω -closed if every decreasing sequence $a_1 > a_2 > \dots$ in P has a lower bound in P .

Lemma 9.2. Let P be any notion of forcing, let G be a generic filter of P , and suppose $\mathbb{M} \models$ “ P is ω -closed”. Also let $A \in \mathbb{M}$. Then any function from \mathbb{N} to A in $\mathbb{M}[G]$ is already in \mathbb{M} .

Proof. Suppose there is a function $f : \mathbb{N} \rightarrow A$ that is in $\mathbb{M}[G]$ but not in \mathbb{M} . Let τ be a name for f ; by Theorem 7.2 (a), some $p \in G$ forces the statement “ τ is a function from $\check{\mathbb{N}}$ to \check{A} ”. Let X be the set of functions from \mathbb{N} to A in \mathbb{M} . Then $f \notin X$, so again by Theorem 7.2 (b) some $q \in G$ forces $\neg(\tau \in \check{X})$. We may assume $q \leq p$.

Define

$$S = \{ \langle r, n, a \rangle : r \in P, n \in \mathbb{N}, a \in A, \text{ and } r \Vdash \langle \check{n}, \check{a} \rangle \in \tau \}.$$

By Theorem 7.2 (b) the set S belongs to M . We claim that for any $n \in \mathbb{N}$ the set of $r \in P$ such that $\langle r, n, a \rangle \in S$ for some a is dense below p . To see this, let $r' \leq p$ and let G' be a generic filter that contains r' . Since $r' \leq p$,

$$M[G'] \models \text{"val}_{G'}(\tau) \text{ is a function from } \mathbb{N} \text{ to } A\text{"}.$$

Therefore

$$M[G'] \models \text{"}\langle n, a \rangle \in \text{val}_{G'}(\tau)\text{"}$$

for some $a \in A$, and hence some $r \in G'$ forces $\langle \check{n}, \check{a} \rangle \in \tau$. Since $r' \in G'$ we may assume that $r \leq r'$, so the claim is proven.

Now, working in M , choose a sequence (p_n) in P and a sequence (a_n) in A such that $q \geq p_0 \geq p_1 \geq \dots$ and p_n forces $\langle \check{n}, \check{a}_n \rangle \in \tau$. We can do this by the claim. Finally, since P is ω -closed we can find $p' \in P$ less than the entire sequence (p_n) . This is a contradiction, because p' forces both $\langle \check{n}, \check{a}_n \rangle \in \tau$ for all n (since it lies below every p_n) and $\neg(\tau \in \check{X})$ (since it lies below q), yet the function $n \mapsto a_n$ evidently is in M (we already constructed the sequence (a_n) in M), and hence it is in X . This proves the lemma. \blacksquare

Theorem 9.3. *Let P be the notion of forcing defined in M as the set of all countable partial bijections between $\mathcal{P}(\mathbb{N})$ and \aleph_1 , ordered by reverse inclusion. Then the continuum hypothesis holds in $M[G]$, for any generic filter G of P .*

Proof. Let G be a generic filter of P . First we show that $\mathcal{P}(\mathbb{N})^{M[G]} = \mathcal{P}(\mathbb{N})^M$ and $\aleph_1^{M[G]} = \aleph_1^M$. The first statement is equivalent to saying that passing from M to $M[G]$ does not introduce any new functions from \mathbb{N} to $\{0, 1\}$, and the second statement is equivalent to saying that passing from M to $M[G]$ does not introduce a surjection from \mathbb{N} (which is absolute) onto \aleph_1^M . Both assertions follow from the lemma, provided we can show that " P is ω -closed" is true in M . But this is trivial: given any decreasing sequence of functions in P that belongs to M , their union is a function in M which lies below the entire sequence.

Now we must show that $M[G]$ contains a bijection between $\mathcal{P}(\mathbb{N})^{M[G]} = \mathcal{P}(\mathbb{N})^M$ and $\aleph_1^{M[G]} = \aleph_1^M$. By Proposition 6.6, $G \in M[G]$. The elements of G are a directed set of partial bijections, so the union of all functions in G is a bijection between a subset of $\mathcal{P}(\mathbb{N})^M$ and a subset of \aleph_1^M . Let f be this union; it belongs to $M[G]$ since the latter satisfies the union axiom. Moreover, for each $x \in \mathcal{P}(\mathbb{N})^M$ the set of functions in P whose domain contains x is dense, and for each $y \in \aleph_1^M$ the set of functions in P whose image contains y is dense (exercise). Since both of these sets lie in M and G is generic, it follows that f is a bijection between $\mathcal{P}(\mathbb{N})^M$ and \aleph_1^M . Thus the continuum hypothesis is true in $M[G]$. \blacksquare

Exercises

- Let τ and σ be P -names. Find a P -name π such that $\text{val}_G(\pi) = \langle \text{val}_G(\tau), \text{val}_G(\sigma) \rangle$ for any generic filter G . (By definition, $\langle x, y \rangle = \{\{x\}, \{x, y\}\}$ for any sets x and y .)
- Let P be the notion of forcing used in Theorem 9.3. Prove that for each $x \in \mathcal{P}(\mathbb{N})^M$ the set of $f \in P$ whose domain contains x is dense, and for each $y \in \aleph_1^M$ the set of $f \in P$ whose image contains y is dense.
- Exhibit a notion of forcing P such that for any generic filter G of P we have $\aleph_1^M = \aleph_1^{M[G]}$, but $\aleph_2^M < \aleph_2^{M[G]}$.

10. Forcing \neg CH

In the last section we had to force 2^{\aleph_0} , which could have been larger than \aleph_1 in M , to equal \aleph_1 in $M[G]$. This was accomplished using a notion of forcing which told us how to build a bijection between $\mathcal{P}(\mathbb{N})$ and \aleph_1 . In this section we want to force $2^{\aleph_0} > \aleph_1$, and we will do this using a notion of forcing which tells us how to build a set of \aleph_2 distinct functions from \mathbb{N} into $\{0, 1\}$.

The appropriate notion of forcing P is the set of all functions from finite subsets of $\mathbb{N} \times \aleph_2^M$ into $\{0, 1\}$. We call elements of P finite partial functions from $\mathbb{N} \times \aleph_2^M$ into $\{0, 1\}$. Note that we do not have to say "functions in M " because for any $A, B \in M$, any function from a finite subset of A into B can be explicitly listed and hence belongs to M . We order P by setting $f \leq g$ if $\text{dom}(g) \subseteq \text{dom}(f)$ and $f|_{\text{dom}(g)} = g$. Again,

f extends g as an element of P if and only if f extends g as a function, and here f and g are compatible as elements of P if and only if they are compatible as functions.

It will be easy to turn a generic filter of P into a function from $\mathbb{N} \times \aleph_2^M$ into $\{0, 1\}$, and also easy to check that this gives us \aleph_2^M distinct functions from \mathbb{N} into $\{0, 1\}$ in $M[G]$. The crucial problem is to show that $\aleph_2^{M[G]} = \aleph_2^M$. We do this using the following notion.

Definition 10.1. A down-antichain in a notion of forcing P is a set $A \subset P$ such that any two distinct elements of A are incompatible. P is c.c.c. if every down-antichain in P is countable.

Definition 10.2. Let P be a notion of forcing. We say that P preserves cardinals if for any generic filter G of P and any ordinal $\alpha \in M$,

$$M \models \text{“}\alpha \text{ is a cardinal”} \quad \text{implies} \quad M[G] \models \text{“}\alpha \text{ is a cardinal”}.$$

Note that the reverse implication is automatic: if α is not a cardinal in M , then there is a bijection in M between α and some smaller ordinal. This bijection will still exist in $M[G]$, so α cannot become a cardinal there. Thus, to say that P preserves cardinals is to say that the cardinals in $M[G]$ are the same as the cardinals in M . (We already saw that M and $M[G]$ have the same ordinals in § 6, Exercise (d).)

We will now show that every c.c.c. notion of forcing preserves cardinals.

Lemma 10.3. *Let P be a notion of forcing such that $M \models \text{“}P \text{ is c.c.c.} \text{”}$ and let G be a generic filter. Suppose $f : A \rightarrow B$ is a function in $M[G]$ with $A, B \in M$. Then there is a map $F : A \rightarrow \mathcal{P}(B)$ in M such that for each $a \in A$ we have $f(a) \in F(a)$ and $M \models \text{“}F(a) \text{ is countable”}$.*

Proof. Let τ be a P -name for f and let $p \in G$ force “ τ is a function from \check{A} to \check{B} ”. For $a \in A$ define

$$F(a) = \{b \in B : \text{some } q \leq p \text{ forces } \langle \check{a}, \check{b} \rangle \in \tau\}.$$

Then F is in M and since some $q \in G$, and hence some $q \leq p$, must force $\langle \check{a}, \check{b} \rangle \in \tau$ with $b = f(a)$, we have $f(a) \in F(a)$. We need only show that $M \models \text{“}F(a) \text{ is countable”}$.

To see this, fix $a \in A$ and, working in M , for each $b \in F(a)$ find $q_b \leq p$ such that $q_b \Vdash \langle \check{a}, \check{b} \rangle \in \tau$. Then the q_b are incompatible since they force different values of $\text{val}(\tau)(a)$ (and they all lie below p , which forces $\text{val}(\tau)$ to be a function). Since “ P is c.c.c.” is true in M , “ $F(a)$ is countable” must also be true in M . ■

Theorem 10.4. *Let P be a notion of forcing such that $M \models \text{“}P \text{ is c.c.c.} \text{”}$. Then P preserves cardinals.*

Proof. Let G be a generic filter of P and suppose $\alpha \in M$ is a cardinal in M but not in $M[G]$. Then there is a smaller cardinal β in M and a function f in $M[G]$ from β onto α . By the lemma there is then a map $F : \beta \rightarrow \mathcal{P}(\alpha)$ in M such that for each $x \in \beta$ we have $f(x) \in F(x)$ and $M \models \text{“}F(x) \text{ is countable”}$. But then

$$M \models \text{“}\text{card}\left(\bigcup_{x \in \beta} F(x)\right) \leq \aleph_0 \cdot \beta = \beta \text{”}$$

(it is clear that β cannot be finite). However, f is surjective and its image is contained in $\bigcup_{x \in \beta} F(x)$, so that the latter must equal α . We have shown that $M \models \text{“}\text{card}(\alpha) \leq \beta \text{”}$, a contradiction. ■

The main remaining step is to show that the notion of forcing described at the beginning of this section is c.c.c. in M . We do this using the following combinatorial lemma, the Δ -systems lemma. This is provable in ZFC.

Lemma 10.5. *Let A be an uncountable family of finite sets. Then there is an uncountable subfamily $B \subset A$ and a set r such that $a \cap b = r$ for any distinct $a, b \in B$.*

Proof. For each $n \in \mathbb{N}$ let $A_n = \{a \in A : \text{card}(a) = n\}$. Then some A_n is uncountable; fix such a value of n . Let r be a maximal set with the property that $r \subseteq a$ for uncountably many $a \in A_n$, and let $A'_n = \{a \in A_n : r \subseteq a\}$. Then A'_n is uncountable and any $x \notin r$ belongs to only countably many $a \in A'_n$. Finally, by Zorn's lemma let B be a maximal subset of A'_n with the property that $a \cap b = r$ for any distinct $a, b \in B$. If B were countable then $\bigcup_{a \in B} a$ would be countable and each x in $(\bigcup_{a \in B} a) - r$ would belong to only countably many $a \in A'_n$, so there would exist an $a \in A'_n$ whose intersection with every $b \in B$ is r , contradicting maximality of B . Thus B must be uncountable. ■

Theorem 10.6. *Let P be the notion of forcing consisting of all finite partial functions from $\mathbb{N} \times \aleph_2^M$ into $\{0, 1\}$, ordered by reverse inclusion. Then the continuum hypothesis fails in $M[G]$, for any generic filter G of P .*

Proof. Let G be a generic filter of P . Working in M , we show that P is c.c.c. Let A be an uncountable subset of P . Applying Lemma 10.5 (which is provable in ZFC, hence true in M) to the domains of the functions in A , we infer that there is an uncountable subset B of A and a (finite) set $r \subseteq \mathbb{N} \times \aleph_2^M$ such that $\text{dom}(f) \cap \text{dom}(g) = r$ for any distinct $f, g \in B$. But there are only finitely many (specifically, $2^{\text{card}(r)}$) possible choices for $f|_r$, so there must be some distinct $f, g \in B$ with $f|_r = g|_r$. Since the union of this f and g is a finite partial function, they are compatible. So, working in M , we have shown that any uncountable subset of P contains compatible functions, and hence that P is c.c.c.

It now follows from Theorem 10.4 that $\aleph_2^{M[G]} = \aleph_2^M$. By Proposition 6.6, $G \in M[G]$. Any two elements of G are compatible partial functions from $\mathbb{N} \times \aleph_2^M$ into $\{0, 1\}$, so the union of all functions in G is a function from a subset of $\mathbb{N} \times \aleph_2^M$ into $\{0, 1\}$. Let F be this union; it belongs to $M[G]$. Moreover, for each $\langle n, \alpha \rangle \in \mathbb{N} \times \aleph_2^M$ the set of finite partial functions whose domain contains $\langle n, \alpha \rangle$ is dense and lies in M , so G must intersect this set and hence F is a function from $\mathbb{N} \times \aleph_2^M$ into $\{0, 1\}$.

Working in $M[G]$, for each $\alpha < \aleph_2$ let

$$S_\alpha = \{n \in \mathbb{N} : F(n, \alpha) = 1\}.$$

For any distinct $\alpha, \beta < \aleph_2^M$ the set

$$\{f \in P : (\exists n \in \mathbb{N})(\langle n, \alpha \rangle, \langle n, \beta \rangle \in \text{dom}(f) \text{ and } f(n, \alpha) \neq f(n, \beta))\}$$

is dense and belongs to M . Since G is generic, this shows that $S_\alpha \neq S_\beta$. Thus

$$M[G] \models \text{“}\{S_\alpha : \alpha < \aleph_2^M\} \text{ is a family of distinct subsets of } \mathbb{N} \text{ of cardinality } \aleph_2^M\text{”},$$

which shows that the continuum hypothesis fails in $M[G]$. ■

Exercises

- For any $A, B \in M$ prove that (1) the set of all finite subsets of A and (2) the set of all functions from a finite subset of A into B are absolute.
- Exhibit a notion of forcing P such that for any generic filter G , $M[G] \models \text{“}\aleph_1^M \text{ is countable”}$.
- Prove that it is relatively consistent with ZFC that $2^{\aleph_0} \geq \aleph_{\omega_1}$.

11. Application: families of analytic functions

We work in ZFC. The following theorem is due to Paul Erdős.

Theorem 11.1. *The continuum hypothesis is true if and only if there is an uncountable family of analytic functions f_α on the complex plane such that for each $z \in \mathbb{C}$ the set of values $\{f_\alpha(z)\}$ is countable.*

Proof. (\Leftarrow) Suppose CH fails. Let $\{f_\alpha : \alpha < \aleph_1\}$ be an uncountable family of distinct analytic functions on the complex plane; we will find a point in \mathbb{C} at which the family takes on uncountably many values. (It is sufficient to consider the case that the family has cardinality \aleph_1 since we can always discard extra functions.)

We invoke a theorem from complex analysis which states that if f and g are analytic functions and the set of points at which they agree has a cluster point, then $f = g$. It follows that two distinct analytic functions can agree on at most countably many points in \mathbf{C} . (If they agreed on uncountably many points, then for some n the compact disk $\{|z| \leq n\}$ would contain infinitely many points of agreement, and there would be a cluster point in this disk.)

For each distinct $\alpha, \beta < \aleph_1$, let

$$S(\alpha, \beta) = \{z \in \mathbf{C} : f_\alpha(z) = f_\beta(z)\}.$$

Then $\bigcup_{\alpha \neq \beta} S(\alpha, \beta)$ is a union of $\aleph_1^2 = \aleph_1$ countable sets and hence has size at most \aleph_1 . Since $\text{card}(\mathbf{C}) = 2^{\aleph_0}$ and we are assuming $2^{\aleph_0} > \aleph_1$, it follows that there exists a point $z_0 \in \mathbf{C}$ that does not belong to any $S(\alpha, \beta)$. This means that $\alpha \neq \beta$ implies $f_\alpha(z_0) \neq f_\beta(z_0)$, so that the family of functions $\{f_\alpha\}$ takes on uncountably many values at the point z_0 .

(\Rightarrow) Suppose CH holds. Let S be the set of complex numbers $a+bi$ with a and b both rational. Since $2^{\aleph_0} = \aleph_1$, we can enumerate the complex plane as $\mathbf{C} = \{z_\alpha : \alpha < \aleph_1\}$. We will construct a family $\{f_\beta : \beta < \aleph_1\}$ of distinct analytic functions such that $f_\beta(z_\alpha) \in S$ for all $\beta > \alpha$. This implies that the set of values of the f_β on each z_α is countable, since (fixing α) $\{f_\beta(z_\alpha) : \beta \leq \alpha\}$ is trivially countable and $\{f_\beta(z_\alpha) : \beta > \alpha\} \subseteq S$ is also countable.

We construct the f_β by transfinite recursion. Suppose f_β has been defined for all $\beta < \gamma$. The set $\{f_\beta : \beta < \gamma\}$ is countable, so we can reorder it as $\{g_n : n \in \mathbf{N}\}$. Similarly, we can reorder $\{z_\alpha : \alpha < \gamma\}$ as $\{w_n : n \in \mathbf{N}\}$. We will construct f_γ so as to satisfy

- (i) $f_\gamma(w_n) \in S$ for all n and
- (ii) $f_\gamma(w_n) \neq g_n(w_n)$ for all n .

It will follow from (i) that $f_\gamma(z_\alpha) \in S$ for all $\alpha < \gamma$, and from (ii) that $f_\gamma \neq f_\beta$ for all $\beta < \gamma$.

We let f_γ be a function of the form

$$f_\gamma(z) = \epsilon_0 + \sum_{n=1}^{\infty} \epsilon_n \prod_{i=0}^{n-1} (z - w_i)$$

where ϵ_n is chosen small enough so that $|\epsilon_n \prod_{i=0}^{n-1} (z - w_i)| \leq 2^{-n}$ on $\{|z| \leq n\}$. This implies that the sum converges uniformly on compact sets, so that it defines an analytic function. We can choose the ϵ_n sequentially, and arrange both conditions (i) and (ii) for w_n when choosing ϵ_n since future terms of the sum vanish on w_n . ■

Exercises

- (a) Let $\{f_n : n \in \mathbf{N}\}$ be infinitely many distinct analytic functions on \mathbf{C} . Prove that there exists $z \in \mathbf{C}$ such that the set of values $\{f_n(z) : n \in \mathbf{N}\}$ is infinite.
- (b) Prove that there are 2^{\aleph_0} distinct continuous functions from \mathbf{C} into $[0, 1]$ which collectively take at most two values at each point. (Hint: do this first for functions from \mathbf{R} into $[0, 1]$.) Can this be done for infinitely differentiable functions?
- (c) Prove that for any uncountable family of complex polynomials, there exists a point in \mathbf{C} at which the family takes on uncountably many values.

References

P. Erdős, An interpolation problem associated with the continuum hypothesis, *Michigan Math. J.* 11 (1964), 9-10.

12. Application: self-homeomorphisms of $\beta\mathbb{N} - \mathbb{N}$, I

We work in ZFC. In this section we show that CH implies there are nontrivial homeomorphisms from the topological space $\beta\mathbb{N} - \mathbb{N}$ onto itself. What we mean by “nontrivial” is explained below. It is also relatively consistent with ZFC that all self-homeomorphisms of $\beta\mathbb{N} - \mathbb{N}$ are trivial; we will discuss this later.

Definition 12.1. A filter over a set X is a family \mathcal{F} of subsets of X which is closed under enlargement ($A \in \mathcal{F}$ and $A \subseteq B$ implies $B \in \mathcal{F}$) and finite intersections ($A, B \in \mathcal{F}$ implies $A \cap B \in \mathcal{F}$). (That is, it is a filter of the poset $\mathcal{P}(X)$ in the sense of Definition 6.1 (c).) An ultrafilter is a maximal proper filter.

We consider ultrafilters over \mathbb{N} . For any $n \in \mathbb{N}$ there is a trivial fixed ultrafilter $\mathcal{U}_{(n)} = \{A \subseteq \mathbb{N} : n \in A\}$. However, there are other ultrafilters; the cofinite filter $\mathcal{F}_{cof} = \{A \subseteq \mathbb{N} : \mathbb{N} - A \text{ is finite}\}$ is a filter, and by Zorn’s lemma, every filter is contained in an ultrafilter, but no ultrafilter containing \mathcal{F}_{cof} can be fixed. An ultrafilter which is not fixed is called free.

Proposition 12.2. *A filter \mathcal{F} over \mathbb{N} is an ultrafilter if and only if for every $A \subseteq \mathbb{N}$, either $A \in \mathcal{F}$ or $\mathbb{N} - A \in \mathcal{F}$ (but not both).*

Proof. Exercise. ■

Definition 12.3. The Stone-Čech compactification of \mathbb{N} , $\beta\mathbb{N}$, is the set of all ultrafilters over \mathbb{N} with the topology whose basic open sets are, for every $A \subseteq \mathbb{N}$, the set U_A of all ultrafilters containing A .

We record basic information about $\beta\mathbb{N}$ in the next result.

Proposition 12.4. *The sets U_A form a basis for a topology on $\beta\mathbb{N}$. This topology is compact, Hausdorff, and totally disconnected. The map $n \mapsto \mathcal{U}_{(n)}$ embeds \mathbb{N} into $\beta\mathbb{N}$ as a dense open set of isolated points.*

Proof. Exercise. ■

It follows from Proposition 12.4 that $\beta\mathbb{N} - \mathbb{N}$ — the subspace of $\beta\mathbb{N}$ consisting of all free ultrafilters — is also compact, Hausdorff, and totally disconnected. It has a basis consisting of the sets

$$\bar{U}_A = U_A \cap (\beta\mathbb{N} - \mathbb{N}) = \{\text{free ultrafilters which contain } A\}$$

for all $A \subseteq \mathbb{N}$. We now introduce the “obvious” self-homeomorphisms of $\beta\mathbb{N} - \mathbb{N}$.

Definition 12.5. An almost permutation of \mathbb{N} is a bijection between two cofinite subsets of \mathbb{N} .

Proposition 12.6. *Let $\phi : S \cong T$ be an almost permutation of \mathbb{N} . For any free ultrafilter \mathcal{U} over \mathbb{N} define $\tilde{\phi}(\mathcal{U})$ to be*

$$\tilde{\phi}(\mathcal{U}) = \{\phi(A \cap S) \cup B : A \in \mathcal{U} \text{ and } B \subseteq \mathbb{N} - T\}$$

(where $\phi(A \cap S) = \{\phi(n) : n \in A \cap S\}$). Then $\tilde{\phi}$ is a self-homeomorphism of $\beta\mathbb{N} - \mathbb{N}$.

Proof. If $S = T = \mathbb{N}$ the proposition is trivial. We need only consider the case when $S = \mathbb{N}$, $T = \mathbb{N} - \{0\}$, and ϕ is the unilateral shift $\phi(n) = n + 1$; this is sufficient because every almost permutation of \mathbb{N} can be expressed as a composition of permutations of \mathbb{N} and some power of the unilateral shift or its inverse. (Apply a permutation to map S to $\{j, j + 1, \dots\}$ where $j = \text{card}(\mathbb{N} - S)$, then compose with ϕ^{k-j} to take this set to $\{k, k + 1, \dots\}$ where $k = \text{card}(\mathbb{N} - T)$, and then compose with another permutation to take $\{k, k + 1, \dots\}$ onto T in a way that matches the given almost permutation.)

Suppose now that ϕ is the unilateral shift. The verification that $\tilde{\phi}$ is a well-defined bijection of $\beta\mathbb{N} - \mathbb{N}$ with itself is routine. It is a homeomorphism because it permutes the basic open sets \bar{U}_A . Specifically, $\tilde{\phi}$ takes \bar{U}_A to $\bar{U}_{A'}$ where $A' = \{n + 1 : n \in A\}$. Every basic open set is of the form $\bar{U}_{A'}$ for some A because in general $\bar{U}_A = \bar{U}_B$ provided the symmetric difference of A and B is finite. ■

We call a self-homeomorphism of $\beta\mathbb{N} - \mathbb{N}$ arising from an almost permutation of \mathbb{N} trivial. We will show that CH implies there are nontrivial self-homeomorphisms of $\beta\mathbb{N} - \mathbb{N}$. This is done using the poset $\mathcal{P}(\mathbb{N})/fin$ defined as follows.

Definition 12.7. Let $\mathcal{P}(\mathbb{N})/fin$ be the quotient of $\mathcal{P}(\mathbb{N})$ by the equivalence relation that makes two subsets of \mathbb{N} equivalent if their symmetric difference is finite. The natural order relation on $\mathcal{P}(\mathbb{N})/fin$ is: the class of A less than the class of B if A is essentially contained in B , i.e., all but finitely many elements of A are also in B .

Lemma 12.8. *The poset of clopen subsets of $\beta\mathbb{N} - \mathbb{N}$ under containment is order-isomorphic to $\mathcal{P}(\mathbb{N})/fin$. This induces a 1-1 correspondence between the self-homeomorphisms of $\beta\mathbb{N} - \mathbb{N}$ and the order-automorphisms of $\mathcal{P}(\mathbb{N})/fin$.*

Proof. First, we claim that the clopen subsets of $\beta\mathbb{N} - \mathbb{N}$ are precisely the basic open sets \bar{U}_A for $A \subseteq \mathbb{N}$. It follows from Proposition 12.2 that the complement of \bar{U}_A is $\bar{U}_{\mathbb{N}-A}$, so the basic open sets are all clopen. Now any open subset can be expressed as a union of basic open sets, and since $\beta\mathbb{N} - \mathbb{N}$ is compact any closed subset is compact, so any clopen subset is a union of finitely many basic open sets. However, the union $\bar{U}_{A_1} \cup \dots \cup \bar{U}_{A_n}$ equals the single basic open set $\bar{U}_{A_1 \cup \dots \cup A_n}$. So every clopen set is a basic open set. This completes the proof of the claim.

If the symmetric difference of A and B is finite then $\bar{U}_A = \bar{U}_B$, because any free ultrafilter contains A if and only if it contains B . However, if the symmetric difference of A and B is infinite then there is a free ultrafilter that contains one but not the other: say $A - B$ is infinite and extend the filter generated by $A - B$ and the cofinite subsets of \mathbb{N} to an ultrafilter. So the map taking \bar{U}_A to the class of A is a bijection between the clopen subsets of $\beta\mathbb{N} - \mathbb{N}$ and the elements of $\mathcal{P}(\mathbb{N})/fin$. We have that A is essentially contained in B if and only if every free ultrafilter that contains A also contains B if and only if $\bar{U}_A \subseteq \bar{U}_B$. So the bijection is an order-isomorphism. This proves the first assertion.

It is clear that any self-homeomorphism of $\beta\mathbb{N} - \mathbb{N}$ induces an order-automorphism of the poset of clopen subsets of $\beta\mathbb{N} - \mathbb{N}$. Conversely, every order-automorphism of the poset of clopen subsets permutes the maximal proper filters of this poset. But by compactness of $\beta\mathbb{N} - \mathbb{N}$, the intersection of any proper filter of clopen sets is nonempty, and if x is any point in the intersection then the original filter must be contained in the filter consisting of all clopen sets that contain x . This shows that the maximal proper filters are precisely the filters of the form: all clopen sets containing the point x , as x ranges over $\beta\mathbb{N} - \mathbb{N}$. We conclude that any order-automorphism of the poset of clopen subsets permutes the points of $\beta\mathbb{N} - \mathbb{N}$, and since this permutation takes basic open sets to basic open sets it is a self-homeomorphism. We have established that any self-homeomorphism of $\beta\mathbb{N} - \mathbb{N}$ induces an order-automorphism of $\mathcal{P}(\mathbb{N})/fin$, and conversely; finally, it is routine to check that these maps are inverse to each other. \blacksquare

Lemma 12.9. *Let (A_n) , (B_n) , and (C_n) be sequences of subsets of \mathbb{N} such that every A_i is essentially contained in every B_j , but not vice versa; no C_k is essentially contained in any A_i ; and no C_k essentially contains any B_j . Also suppose that the A_n are directed upward and the B_n are directed downward under almost containment. Then there are two subsets A and B of \mathbb{N} with infinite symmetric difference which essentially contain (but are not essentially contained in) each A_n , are essentially contained in (but do not essentially contain) each B_n , and neither essentially contain nor are essentially contained in any C_n .*

Proof. We construct sequences

$$S_1 \subset S'_1 \subset S_2 \subset S'_2 \subset \dots$$

and

$$T_1 \subset T'_1 \subset T_2 \subset T'_2 \subset \dots$$

such that $S_n \cap T_n = \emptyset$ for all n . We think of the elements of S_n as required and the elements of T_n as forbidden. Start with $S_1 = T_1 = \emptyset$. Given S_n and T_n , construct S'_n from S_n by adding A_n , then subtracting T_n , then adding a point in $\mathbb{N} - (A_i \cup T_n)$ for each $1 \leq i \leq n$. Then construct T'_n by adding $\mathbb{N} - B_n$, subtracting S'_n , and adding a point in $B_i - S'_n$ for each $1 \leq i \leq n$. Then construct S_{n+1} by adding to S'_n

one point in $\mathbb{N} - (C_i \cup T'_n)$ for each $1 \leq i \leq n$ such that S'_n is essentially contained in C_i , and construct T_{n+1} by adding to T'_n one point in $C_i - S_{n+1}$ for each $1 \leq i \leq n$ such that $\mathbb{N} - T'_n$ essentially contains C_i . The construction can proceed because we inductively have that S_n has finite symmetric difference with $A_1 \cup \dots \cup A_n$ and T_n has finite symmetric difference with $\mathbb{N} - (B_1 \cap \dots \cap B_n)$. Finally, let $A = \bigcup S_n$. It is immediate that A essentially contains every A_n , and by the construction of S'_n it contains infinitely many points outside of each A_n . Since A does not intersect $\bigcup T_n$, it is essentially contained in each B_n , and by the construction of T'_n it avoids infinitely many points in each B_n . By the construction of S_n and T_n we have that A neither essentially contains nor is essentially contained in any C_n . To find a set B with the same properties that essentially contains but is not essentially contained in A , add A to the family $\{A_n\}$ and apply the construction again. \blacksquare

Theorem 12.10. *Assume the continuum hypothesis. Then there exist nontrivial self-homeomorphisms of $\beta\mathbb{N} - \mathbb{N}$.*

Proof. By Lemma 12.8, it suffices to find a nontrivial order-automorphism of $\mathcal{P}(\mathbb{N})/fin$. There are 2^{\aleph_0} trivial order-automorphisms; enumerate them as $\{\psi_\alpha : \alpha < \aleph_1\}$. The cardinality of $\mathcal{P}(\mathbb{N})/fin$ is also 2^{\aleph_0} , so enumerate its elements as $\{x_\alpha : \alpha < \aleph_1\}$. We recursively construct countable subsets \mathcal{A}_α of $\mathcal{P}(\mathbb{N})/fin$ which are closed under finite unions, finite intersections, and complements, together with injective maps $\phi_\alpha : \mathcal{A}_\alpha \rightarrow \mathcal{P}(\mathbb{N})/fin$, such that, with $\mathcal{B}_\alpha = \phi_\alpha(\mathcal{A}_\alpha)$,

- (i) if $\beta \leq \alpha$ then $\mathcal{A}_\beta \subseteq \mathcal{A}_\alpha$, $\mathcal{B}_\beta \subseteq \mathcal{B}_\alpha$, and $\phi_\alpha|_{\mathcal{A}_\beta} = \phi_\beta$;
- (ii) ϕ_α preserves finite unions, finite intersections, and complements;
- (iii) ϕ_α disagrees with ψ_α on \mathcal{A}_α ; and
- (iv) x_α is in both \mathcal{A}_α and \mathcal{B}_α .

At stage α of the construction, if $x_\alpha \notin \bigcup_{\beta < \alpha} \mathcal{A}_\beta$ then we let $y_\alpha = x_\alpha$; otherwise we choose y_α arbitrarily in the complement of $\bigcup_{\beta < \alpha} \mathcal{A}_\beta$. Then we let $\tilde{\mathcal{A}}_\alpha$ be generated by $\bigcup_{\beta < \alpha} \mathcal{A}_\beta$ and y_α under the operations of union, intersection, and complement, extend the previously defined ϕ_β to an embedding $\tilde{\phi}_\alpha$ of $\tilde{\mathcal{A}}_\alpha$ in $\mathcal{P}(\mathbb{N})/fin$, and let $\tilde{\mathcal{B}}_\alpha = \tilde{\phi}_\alpha(\tilde{\mathcal{A}}_\alpha)$. If $\tilde{\mathcal{B}}_\alpha$ contains x_α then we let $\mathcal{A}_\alpha = \tilde{\mathcal{A}}_\alpha$, $\mathcal{B}_\alpha = \tilde{\mathcal{B}}_\alpha$, and $\phi_\alpha = \tilde{\phi}_\alpha$; otherwise we let \mathcal{B}_α be generated by $\tilde{\mathcal{B}}_\alpha$ and x_α under the operations of union, intersection, and complement, extend $\tilde{\phi}_\alpha^{-1}$ to \mathcal{B}_α , and let \mathcal{A}_α be the image of \mathcal{B}_α under the extension and ϕ_α the inverse of the extension.

The point is that there are always at least two ways to define $\tilde{\phi}_\alpha$, hence at most one of them can agree with ψ_α , and we can ensure condition (iii) by choosing the other one. To define the extension, partition $\bigcup_{\beta < \alpha} \mathcal{A}_\beta$ into those elements which lie below y_α , those which lie above y_α , and those which are incomparable to y_α . This corresponds to a partition of $\bigcup_{\beta < \alpha} \mathcal{B}_\beta$, which can be lifted to three countable families of subsets of \mathbb{N} that satisfy the hypotheses of Lemma 12.9. That lemma then provides us with two suitable targets for y_α ; this determines the construction of $\tilde{\phi}_\alpha$, and it is straightforward to check that this does produce a well-defined map with the desired properties.

Finally, we define an order-automorphism ϕ of $\mathcal{P}(\mathbb{N})/fin$ by taking the union of the maps ϕ_α . It is injective since each ϕ_α is injective, it is surjective since for each α its image contains $x_\alpha \in \mathcal{B}_\alpha$, and it is an order-isomorphism since each $\phi_\alpha : \mathcal{A}_\alpha \cong \mathcal{B}_\alpha$ is an order-isomorphism. It disagrees with each ψ_α by construction, so it is a nontrivial order-automorphism. \blacksquare

Exercises

- (a) Prove Proposition 12.2.
- (b) Prove Proposition 12.4. (Hint: for compactness, first look at Exercise (c).)
- (c) In the proof of Lemma 12.8 show that $\tilde{U}_{A_1} \cup \dots \cup \tilde{U}_{A_n} = \tilde{U}_{A_1 \cup \dots \cup A_n}$.
- (d) Prove that any self-homeomorphism of $\beta\mathbb{N}$ arises from a permutation of \mathbb{N} .

References

J. van Mill, An introduction to $\beta\omega$, in *Handbook of Set-Theoretic Topology*, pp. 503-567, 1984.

13. Application: pure states on $\mathcal{B}(H)$

We work in ZFC. Let H be a separable infinite-dimensional complex Hilbert space and let $\mathcal{B}(H)$ be the set of all bounded linear operators from H to itself.

Definition 13.1. A (concrete) C*-algebra is a norm-closed linear subspace of $\mathcal{B}(H)$ that is stable under adjoints and products (i.e., composition of operators). It is unital if it contains the identity operator on H . A state on a unital C*-algebra \mathcal{A} is a bounded linear functional $f : \mathcal{A} \rightarrow \mathbb{C}$ such that $\|f\| = f(I) = 1$, where I is the identity operator. It is pure if it cannot be expressed as $f = (f_1 + f_2)/2$ for distinct states f_1 and f_2 .

For example, if v is a unit vector in H then the map $f_v : \mathcal{A} \rightarrow \mathbb{C}$ is a state on $\mathcal{B}(H)$ because

$$\langle Iv, v \rangle = \langle v, v \rangle = 1$$

(so that $f_v(I) = 1$ and consequently $\|f_v\| \geq 1$) and for any operator $A \in \mathcal{B}(H)$

$$|f_v(A)| = |\langle Av, v \rangle| \leq \|Av\| \|v\| \leq \|A\| \|v\|^2 = \|A\|$$

(so that $\|f_v\| \leq 1$). In fact it is pure, though this is a little harder to show. Any convex combination of states is a state; thus the set of all states is convex, and the pure states are its extreme points.

Let \mathcal{U} be an ultrafilter over \mathbb{N} and let (a_n) be a bounded sequence of complex numbers. For each $A \in \mathcal{U}$ let S_A be the closure of $\{a_n : n \in A\}$; then each S_A is compact, and they have the finite intersection property because

$$S_{A_1 \cap \dots \cap A_n} \subseteq S_{A_1} \cap \dots \cap S_{A_n}.$$

So $\bigcap_{A \in \mathcal{U}} S_A$ is nonempty. Furthermore, it must consist of a single point because any two distinct points in \mathbb{C} can be separated by a pair of open sets U and V ; since either $\{n : a_n \in U\}$ or $\{n : a_n \notin U\}$ is in \mathcal{U} , but not both, some S_A is contained in either the closure of U or the complement of U . Thus $\bigcap_{A \in \mathcal{U}} S_A = \{a\}$ for some complex number a . We write $\lim_{\mathcal{U}} a_n = a$. (This construction would work for ultrafilters over any set, a fact we will use in the proof of Lemma 13.4.)

The following theorem is due to Joel Anderson. We omit the proof.

Theorem 13.2. *Let (e_n) be an orthonormal basis of H and let \mathcal{U} be an ultrafilter over \mathbb{N} . Then*

$$f(A) = \lim_{\mathcal{U}} \langle Ae_n, e_n \rangle$$

defines a pure state on $\mathcal{B}(H)$.

The states described in Theorem 13.2 are called diagonalizable, and until recently it was not known whether there are any pure states on $\mathcal{B}(H)$ that are not diagonalizable. Charles Akemann and I recently proved that CH implies there exist non-diagonalizable pure states. The relative consistency with ZFC of the statement “all pure states are diagonalizable” remains open.

We will give the non-diagonalizable construction, leaving out the proof of one key C*-algebraic lemma. An operator in $\mathcal{B}(H)$ is finite rank if its range is finite-dimensional. We omit the proof of the following result.

Lemma 13.3. *Let $\mathcal{A} \subseteq \mathcal{B}(H)$ be a separable unital C*-algebra which contains all finite rank operators. Let f be a pure state on \mathcal{A} that is zero on all finite rank operators and let (e_n) be an orthonormal basis of H . Then there is a subset $S \subseteq \mathbb{N}$ and a pure state g on $\mathcal{B}(H)$ that extends f such that $0 < g(P) < 1$ where P is the operator defined by*

$$Pe_n = \begin{cases} e_n & \text{if } n \in S \\ 0 & \text{if } n \notin S. \end{cases}$$

We call the operator P in Lemma 13.3 an orthogonal projection that is diagonalized by (e_n) . Observe that for any ultrafilter \mathcal{U} over \mathbb{N} the pure state $A \mapsto \lim_{\mathcal{U}} \langle Ae_n, e_n \rangle$ takes either the value 1 or the value 0 on P , depending on whether S or $\mathbb{N} - S$ is in \mathcal{U} .

Lemma 13.4. *Assume the continuum hypothesis. Let f be a pure state on $\mathcal{B}(H)$ and let $\mathcal{A} \subseteq \mathcal{B}(H)$ be a separable unital C^* -algebra. Then there is a separable unital C^* -algebra \mathcal{B} that contains \mathcal{A} such that the restriction of f to \mathcal{B} is pure.*

Proof. The cardinality of $\mathcal{B}(H)$ is 2^{\aleph_0} . By the continuum hypothesis we can enumerate the elements of $\mathcal{B}(H)$ as $\{x_\alpha : \alpha < \aleph_1\}$. For each $\alpha < \aleph_1$ let \mathcal{A}_α be the unital C^* -algebra generated by \mathcal{A} and $\{x_\beta : \beta < \alpha\}$.

Let $x = x_\gamma \in \mathcal{B}(H)$. We claim that for sufficiently large $\alpha > \gamma$ we have, for any states f_1 and f_2 on \mathcal{A}_α ,

$$f|_{\mathcal{A}_\alpha} = (f_1 + f_2)/2 \quad \text{implies} \quad f_1(x) = f_2(x).$$

(If this condition holds we say that $f|_{\mathcal{A}_\alpha}$ is pure on x .) To see this, suppose the claim fails. Then there exist $\epsilon > 0$, an unbounded set $T \subseteq \aleph_1$, and for each $\alpha \in T$ states f_1^α and f_2^α on \mathcal{A}_α such that

$$f|_{\mathcal{A}_\alpha} = (f_1^\alpha + f_2^\alpha)/2 \quad \text{and} \quad |f_1^\alpha(x) - f_2^\alpha(x)| \geq \epsilon. \quad (*)$$

(If no such ϵ and T existed, then for each $n \in \mathbb{N}$ we could find $\beta_n < \aleph_1$ such that $(*)$ cannot be achieved with $\epsilon = 1/n$ for any $\alpha \geq \beta_n$. This would imply that $f|_{\mathcal{A}_\alpha}$ is pure on x for all $\alpha > \sup \beta_n$, contradicting our assumption that the claim fails.)

Let \mathcal{U} be an ultrafilter over T that contains the set $\{\alpha \in T : \alpha > \beta\}$ for each $\beta < \aleph_1$, and define states f_1 and f_2 on $\mathcal{B}(H)$ by $f_1(y) = \lim_{\mathcal{U}} f_1^\alpha(y)$ and $f_2(y) = \lim_{\mathcal{U}} f_2^\alpha(y)$ for all $y \in \mathcal{B}(H)$. Then $f = (f_1 + f_2)/2$ and $|f_1(x) - f_2(x)| \geq \epsilon$, so that $f \neq f_1$. This contradicts the assumption that f is pure, so the claim is established.

Now for any $\alpha < \aleph_1$, we can find a countable dense subset of \mathcal{A}_α , and the claim implies that for sufficiently large β , $f|_{\mathcal{A}_\beta}$ is pure on every element of this subset. Since the subset is dense, $f|_{\mathcal{A}_\beta}$ is pure on every $x \in \mathcal{A}_\beta$. Now construct a sequence (α_n) by setting $\alpha_0 = 0$ and choosing $\alpha_{n+1} > \alpha_n$ such that $f|_{\mathcal{A}_{\alpha_{n+1}}}$ is pure on every $x \in \mathcal{A}_{\alpha_n}$. Letting $\beta = \sup \alpha_n$, we then have that $f|_{\mathcal{A}_\beta}$ is pure on every element of $\bigcup_{\alpha < \beta} \mathcal{A}_\alpha$; since this is a dense subset of \mathcal{A}_β , $f|_{\mathcal{A}_\beta}$ is a pure state on \mathcal{A}_β . Thus $\mathcal{B} = \mathcal{A}_\beta$ is a separable unital C^* -algebra that contains \mathcal{A} such that the restriction of f to \mathcal{B} is pure. \blacksquare

Theorem 13.5. *Assume the continuum hypothesis. Then there is a pure state on $\mathcal{B}(H)$ that is not diagonalizable.*

Proof. There are 2^{\aleph_0} elements of $\mathcal{B}(H)$ and 2^{\aleph_0} orthonormal bases of H . Enumerate these as $\{x_\alpha : \alpha < \aleph_1\}$ and $\{(e_n^\alpha) : \alpha < \aleph_1\}$. We recursively construct a nested transfinite sequence of separable unital C^* -algebras \mathcal{A}_α together with pure states f_α on \mathcal{A}_α such that for all $\alpha < \aleph_1$ we have

- (i) $x_\alpha \in \mathcal{A}_{\alpha+1}$;
- (ii) if $\beta < \alpha$ then $f_\alpha|_{\mathcal{A}_\beta} = f_\beta$; and
- (iii) $\mathcal{A}_{\alpha+1}$ contains an orthogonal projection P_α diagonalized by (e_n^α) such that $0 < f_{\alpha+1}(P_\alpha) < 1$.

Begin by letting \mathcal{A}_0 be the closure of $\{A + zI : A \text{ is finite rank and } z \in \mathbb{C}\}$ and defining $f_0(A + zI) = z$. At successor stages, use Lemma 13.3 to find an orthogonal projection P_α that is diagonalized by (e_n^α) and a pure state g on $\mathcal{B}(H)$ such that $g|_{\mathcal{A}_\alpha} = f_\alpha$ and $0 < g(P_\alpha) < 1$. Then use Lemma 13.4 to find a separable unital C^* -algebra $\mathcal{A}_{\alpha+1}$ which contains \mathcal{A}_α , x_α , and P_α and such that the restriction $f_{\alpha+1}$ of g to $\mathcal{A}_{\alpha+1}$ is pure. At limit ordinals α , let \mathcal{A}_α be the closure of $\bigcup_{\beta < \alpha} \mathcal{A}_\beta$. In this case the state f_α is determined by continuity and the condition that $f_\alpha|_{\mathcal{A}_\beta} = f_\beta$; it is pure because if g_1 and g_2 are states on \mathcal{A}_α such that $f_\alpha = (g_1 + g_2)/2$, then for all $\beta < \alpha$ purity of f_β implies that g_1 and g_2 have the same restriction to \mathcal{A}_β , so that $g_1 = g_2$. This completes the description of the recursive construction.

Finally, define a state f on $\mathcal{B}(H)$ by letting $f|_{\mathcal{A}_\alpha} = f_\alpha$. By the reasoning used immediately above, f is pure, and since $0 < f(P_\alpha) < 1$ for all α , it is not diagonalized by any orthonormal basis (e_n^α) . \blacksquare

Exercises


- (a) Prove that the set of states on a unital C^* -algebra is convex and weak*-compact.
- (b) Let (a_n) and (b_n) be bounded sequences of complex numbers and let \mathcal{U} be an ultrafilter over \mathbb{N} . Prove that $\lim_{\mathcal{U}}(a_n + b_n) = \lim_{\mathcal{U}} a_n + \lim_{\mathcal{U}} b_n$.
- (c) Show that the state f_0 defined on \mathcal{A}_0 in the proof of Theorem 13.5 is pure. (Hint: if P is an orthogonal projection such that $I - P$ has finite rank, then $f_0(P) = 1$. Every finite rank operator is a linear combination of rank one orthogonal projections.)

References

- C. Akemann and N. Weaver, $B(H)$ has a pure state that is not multiplicative on any MASA, *Proc. Nat. Acad. Sci. USA*, to appear.
- J. Anderson, Extreme points in sets of positive linear maps on $B(H)$, *J. Funct. Anal.* 31 (1979), 195-217.

14. The diamond principle

Consider an infinite binary tree. If one vertex is removed from each level except the first, there will still be an infinite path down the tree. Such a path can be recursively constructed starting at the top vertex, since each vertex in the original tree has two immediate successors, so that after vertices are removed at least one immediate successor will always still be available. Thus the construction can proceed indefinitely.



Avoiding forbidden vertices

Now consider the analogous question for the standard \aleph_1 - \aleph_1 -tree. This tree has \aleph_1 levels and each vertex has \aleph_1 immediate successors. The vertices at level α can be labelled by functions from α into \aleph_1 ; if $\alpha < \beta$ then a vertex w at level β lies below a vertex v at level α if the function that labels w extends the function that labels v . Remove one vertex from each level except the first. Will there necessarily still be at least one path all the way down the tree?

The binary tree argument does not work here. The problem is that when we try to recursively construct a path down the tree, the vertex we hit at a limit ordinal is determined by the sequence of vertices chosen up to that point. (If α is a limit ordinal then any function from α to \aleph_1 is determined by its restrictions to all ordinals strictly less than α .) Thus, when attempting to construct a path down the tree, in order to avoid forbidden vertices at limit levels we would have to look ahead somehow. It is not clear that this can be done.

Definition 14.1. A subset C of \aleph_1 is closed if the supremum of every countable subset of C lies in C , and it is unbounded if for every $\alpha < \aleph_1$ there exists $\beta \in C$ such that $\beta > \alpha$. (Thus, C is unbounded if and only if it is uncountable.) A subset of \aleph_1 is stationary if it intersects every closed unbounded subset. The diamond principle (\diamond) asserts that one vertex can be chosen from each level of the standard \aleph_1 - \aleph_1 -tree in such a way that every path down the tree repeatedly meets the chosen vertices on a stationary set of levels.

Less formally: one vertex can be chosen from each level so that every path down the tree repeatedly meets chosen vertices. More formally: there is a sequence $\{h_\alpha : \alpha < \aleph_1\}$ of functions $h_\alpha : \alpha \rightarrow \aleph_1$ such that for any function $f : \aleph_1 \rightarrow \aleph_1$ the set

$$\{\alpha : f|_\alpha = h_\alpha\}$$

is a stationary subset of \aleph_1 .

The diamond principle implies the continuum hypothesis (exercise), so $\neg\Diamond$ follows from $\neg\text{CH}$ and hence is relatively consistent with ZFC. We will now show that \Diamond is also relatively consistent with ZFC. In M , define a notion of forcing P as follows. The elements of P are all sequences $p = \{h_\gamma : \gamma < \alpha\}$ such that α is a countable ordinal and h_γ is a function from γ into \aleph_1 , for each γ . (Thus, p chooses one vertex from each level of the tree up to level α .) We call α the length of p . Order the elements of P by letting $p \leq q$ if the sequence p extends the sequence q .

Theorem 14.2. *Let P be the notion of forcing defined above. Then \Diamond is true in $M[G]$, for any generic filter G .*

Proof. It is clear that $M \models "P \text{ is } \omega\text{-closed}"$, so $\aleph_1^M = \aleph_1^{M[G]}$. As usual, the union of G is a sequence $\{h_\alpha : \alpha < \aleph_1\}$ such that each h_α is a function from α into \aleph_1 . We must show that in $M[G]$ the restrictions of any function $f : \aleph_1^M \rightarrow \aleph_1^M$ agree with the h_α on a stationary set.

Let f be a function in $M[G]$ from \aleph_1^M to itself and let $C \subseteq \aleph_1^M$ be a subset in $M[G]$ such that $M[G] \models "C \text{ is closed and unbounded}"$. Let τ be a P -name for f and let σ be a P -name for C , and find $p \in G$ such that $p \Vdash "\tau \text{ is a function from } \aleph_1 \text{ to itself and } \sigma \text{ is a closed and unbounded subset of } \aleph_1"$.

Working in M , for any $r \leq p$ we now recursively define a sequence $p_0 \geq p_1 \geq \dots$. Let $p_0 = r$. Having chosen p_n , we choose p_{n+1} as follows. Say the length of p_n is α_n . Since $p \Vdash "\sigma \text{ is unbounded}"$, for any generic filter G' that contains p_n we have that $M[G'] \models "there \text{ exists } \beta > \alpha_n \text{ in } \text{val}_{G'}(\sigma)"$. Hence there exists $q_n < p_n$ which forces $\check{\beta}_n \in \sigma$, for some $\beta_n > \alpha_n$. We may assume that the length of q_n is at least β_n . Similarly, we can find $p_{n+1} \leq q_n$ such that p_{n+1} forces " $\tau|_{\check{\alpha}_n} = \check{f}_n$ " for some function $f_n : \alpha_n \rightarrow \aleph_1^M$ in M . (Here we use the fact that P is ω -closed: by Lemma 9.2 and the fact that α_n is countable in M , the evaluation of τ restricted to α_n belongs to M , for any generic filter that contains q_n .)

We have constructed (p_n) in M . Let q be the union of the sequences p_n , so that the length of q is $\alpha^* = \sup \alpha_n$. Also, since $\alpha_n < \beta_n \leq \alpha_{n+1}$ we have $\alpha^* = \sup \beta_n$. Since $q \leq q_n$ for all n , we have $q \Vdash \check{\beta}_n \in \sigma$. Since $q \Vdash "\sigma \text{ is closed}"$ we must have $q \Vdash \check{\alpha}^* \in \sigma$. Finally, let $q' = \{g_\gamma : \gamma < \alpha^* + 1\}$ be the sequence of length $\alpha^* + 1$ which extends q by the one additional function $g_{\alpha^*} : \alpha^* \rightarrow \aleph_1$ defined by taking the union of the functions f_n . Since $q' \leq p_{n+1}$, it forces "the restriction of τ to $\check{\alpha}_n$ equals the restriction of \check{g}_{α^*} to $\check{\alpha}_n$ " for all n ; thus, q' forces "the restriction of τ to $\check{\alpha}^*$ equals \check{g}_{α^*} ".

We conclude that q' forces " $\check{\alpha}^* \in \sigma$ and τ restricted to $\check{\alpha}^*$ is in the union of Γ ". Thus, we have shown that any $r \leq p$ lies above some q' which forces that τ restricted to some element of σ equals a function in the union of Γ . So the set of such q' is dense below p , and since G is generic it must contain such a point. This means that $M[G] \models "(\exists \alpha)(\alpha \in C \text{ and } f|_\alpha = h_\alpha)"$. Thus, in $M[G]$ the set of α for which the restriction of any given function from \aleph_1 to α equals h_α is stationary. ■

Exercises

- Prove that the intersection of countably many closed unbounded subsets of \aleph_1 is a closed unbounded subset of \aleph_1 .
- Prove that \Diamond implies CH.

15. Application: Suslin's problem, I

We work in ZFC. The real line is dense (between any two points there is a third), unbounded (there is no least or greatest element), complete (every bounded set has a least upper bound and a greatest lower bound), and separable. Conversely, it is not too hard to show that any totally ordered set with these properties is order-isomorphic to \mathbb{R} . Suslin's problem asks whether "separable" can be weakened to "there is no uncountable collection of disjoint open intervals" in this characterization. This condition is called c.c.c.

Definition 15.1. A Suslin line is a totally ordered set that is dense, unbounded, complete, and c.c.c. but not order-isomorphic to \mathbb{R} .

The existence of Suslin lines is independent of ZFC. In this section we will show that diamond implies that Suslin lines exist.

The key technical step is a reduction to a problem about trees.

Definition 15.2. A tree is a partially ordered set with the property that for any x in T the set $\{y \in T : y > x\}$ is reverse well-ordered.

(a) If T is a tree and $x \in T$, the height of x is the ordinal reverse isomorphic to $\{y : y > x\}$. A level of T is the set of all vertices of a given height. The height of T is the supremum of the heights of its vertices.

(b) A branch in T is a maximal totally ordered subset, and an antichain is a set of pairwise incomparable elements.

(c) T is a Suslin tree if its height is \aleph_1 , every branch is countable, and every antichain is countable. It is normal if additionally

- (i) it has a unique greatest element;
- (ii) every vertex has infinitely many immediate successors;
- (iii) every vertex has descendants at every level $\alpha < \aleph_1$; and
- (iv) at any limit level, no two distinct vertices have exactly the same predecessors.

Lemma 15.3. *Suslin lines exist if and only if normal Suslin trees exist.*

Proof. (\Rightarrow) Given a Suslin line, we create a Suslin tree as follows. Let I_0 be an arbitrary nondegenerate closed interval in the line. For $\alpha < \aleph_1$ we recursively choose I_α to be a nondegenerate closed interval that does not contain either endpoint of any interval I_β with $\beta < \alpha$. We can do this because the set of all endpoints to be avoided is countable and hence not dense in the line, since a Suslin line cannot be separable. The vertices of the tree are the countable ordinals, ordered by setting $\alpha \preceq \beta$ if $I_\alpha \subseteq I_\beta$.

We verify that the resulting tree is Suslin. The nonexistence of uncountable antichains follows from the fact that one cannot find uncountably many disjoint intervals I_α (since the original line was Suslin). Similarly, any branch corresponds to a nested sequence of intervals; the left endpoints of these intervals are then an increasing sequence in the line, so that they constitute the endpoints of a family of disjoint intervals. Again, since the original line was c.c.c., any such sequence must be countable. So the tree has no uncountable branches.

Finally, it is an exercise to show that every Suslin tree can be converted into a normal Suslin tree.

(\Leftarrow) Given a normal Suslin tree, first assign an order to the immediate successors of each vertex, making them order-isomorphic to the rationals. Then define a Suslin line by letting its points be the branches of the tree, ordered lexicographically according to the orderings just introduced. It is clear that the resulting totally ordered set is dense and has no least or greatest element. It is c.c.c. because for any open interval we can find a vertex x such that all branches passing through x lie in the interval, so an uncountable family of disjoint open intervals would give rise to an uncountable antichain in the original tree. It is not separable because any countable set of branches have heights less than some countable ordinal, and any point at any lower level is contained in an interval of branches that evidently does not meet the original set of branches. Finally, we obtain a Suslin line by completing. \blacksquare

Lemma 15.4. *Let T be a tree of height \aleph_1 and let A be a maximal antichain in T . For each $\alpha < \aleph_1$ let T_α be the set of vertices of height less than α and suppose each T_α is countable. Then*

$$C = \{\alpha : A \cap T_\alpha \text{ is a maximal antichain in } T_\alpha\}$$

is a closed unbounded subset of \aleph_1 .

Proof. For any α , $A \cap T_\alpha$ is an antichain in T_α . If $\alpha = \sup \alpha_n$ and $A \cap T_{\alpha_n}$ is a maximal antichain in each T_{α_n} then any element of T_α belongs to some T_{α_n} and hence is comparable to some element of $A \cap T_\alpha$; thus $\alpha \in C$. So C is closed.

To see that C is unbounded, let $\alpha < \aleph_1$. Then T_α is countable and every element of T_α is comparable to some element of A , so there exists $\alpha_1 \geq \alpha$ such that every element of T_α is comparable to some element of $A \cap T_{\alpha_1}$. Then find $\alpha_2 \geq \alpha_1$ such that every element of T_{α_1} is comparable to some element of $A \cap T_{\alpha_2}$, and so on. Setting $\alpha^* = \sup \alpha_n$, we have that every element of T_{α^*} is comparable to some element of $A \cap T_{\alpha^*}$, i.e., $\alpha^* \in C$. Since $\alpha^* \geq \alpha$ and α was arbitrary, this shows that C is unbounded. ■

Theorem 15.5. *Assume \diamond . Then Suslin lines exist.*

Proof. Using \diamond , fix a sequence $\{h_\alpha : \alpha < \aleph_1\}$ of functions $h_\alpha : \alpha \rightarrow \{0, 1\}$ such that for any function $f : \aleph_1 \rightarrow \{0, 1\}$ the set $\{\alpha : f|_\alpha = h_\alpha\}$ is a stationary subset of \aleph_1 . We construct a normal Suslin tree T whose vertices are all countable ordinals. We take the ordinals in order, i.e., every time we add a vertex it will be the first ordinal not used up to that point.

Let 0 be the top vertex. Having constructed all vertices at level α , we add a countably infinite number of immediate successors to each vertex at level α to obtain the vertices at level $\alpha+1$. We construct the vertices at a limit level α as follows. Let β be the least ordinal not yet used as a vertex and let $S_\beta = \{\gamma < \beta : h_\beta(\gamma) = 0\}$. Using the notation of Lemma 15.4, if S_β is not a maximal antichain in T_α , then for each $\gamma < \beta$ choose a path of height α containing γ (it will be clear inductively that this can always be done) and add a vertex at level α below this path. If S_β is a maximal antichain in T_α , then do the same thing but ensure that the path of height α containing γ also contains a point of S_β . We can do this because S_β is maximal. This completes the description of the construction.

We must show that every branch is countable and every antichain is countable; the other properties of a normal Suslin tree are clear from the construction. Let A be a maximal antichain; by Lemma 15.4 the set of α such that $A \cap T_\alpha$ is maximal in T_α is closed and unbounded. Also, the set of α such that $T_\alpha = \alpha$ is easily seen to be closed and unbounded, so by \diamond there exists α such that $T_\alpha = \alpha$ and $S_\alpha = A \cap T_\alpha$ is a maximal antichain in T_α . By the construction, every vertex at level α then lies below some element of S_α , and this must remain true at all future levels. Thus $A = A \cap T_\alpha$, and hence it is countable. We have shown that every antichain in T is countable.

Finally, there is no uncountable branch because this would easily imply the existence of uncountable antichains. (For each vertex in the branch, choose an immediate successor that does not lie in the branch. The set of these successors is an antichain.) ■

Exercises

(a) Show that if a Suslin tree exists then a normal Suslin tree exists. (Hint: Given a Suslin tree, first remove any vertex that has descendants at only countably many levels. Then remove any vertex that has only one immediate successor, but do not necessarily remove the descendants of such vertices. Then insert vertices to accommodate condition (iv). Next remove all vertices at all successor levels, and finally handle condition (i).)

16. Application: Naimark's problem

We work in ZFC. Recall that C^* -algebras are defined concretely as closed linear subspaces of $B(H)$ that are stable under products and adjoints. We now allow H to be nonseparable. Isomorphic C^* -algebras can be realized as acting on different Hilbert spaces in different ways; a representation of a C^* -algebra \mathcal{A} on a Hilbert space K is a bounded linear map $\pi : \mathcal{A} \rightarrow B(K)$ which preserves products and adjoints. It is irreducible if K cannot be nontrivially decomposed into an orthogonal direct sum $K = K_1 \oplus K_2$ in such a way that each summand is invariant for the action of \mathcal{A} (that is, $\pi(x)K_i \subseteq K_i$ for all $x \in \mathcal{A}$ and $i = 1, 2$). If $\pi' : \mathcal{A} \rightarrow B(K')$ is another representation of \mathcal{A} , we say that π and π' are unitarily equivalent if there is a unitary operator $U : K \rightarrow K'$ such that $\pi(x) = U^{-1}\pi'(x)U$ for all $x \in \mathcal{A}$.

Irreducible representations are related to pure states by the following facts from basic C^* -algebra theory. If π is an irreducible representation of \mathcal{A} on $B(K)$ and v is a unit vector in K , then $x \mapsto \langle \pi(x)v, v \rangle$ is a pure state on \mathcal{A} . Every pure state is realized in this way for some irreducible representation. If \mathcal{A} is unital

then two irreducible representations are unitarily equivalent if and only if (some or any) corresponding pure states are unitarily equivalent in the following sense.

Definition 16.1. Two pure states ρ_1 and ρ_2 on a unital C^* -algebra \mathcal{A} are unitarily equivalent if there is a unitary $u \in \mathcal{A}$ such that $\rho_1(x) = \rho_2(u^*xu)$ for all $x \in \mathcal{A}$.

For any Hilbert space H , the C^* -algebra of compact operators on H , $\mathcal{K}(H)$, is the closure of the set of finite rank operators. Mark Naimark proved that $\mathcal{K}(H)$ has only one irreducible representation up to unitary equivalence, namely the identity representation on H , and he asked whether any other C^* -algebra has only one irreducible representation. Alex Rosenberg quickly showed that there were no separable examples. Recently Charles Akemann and I used diamond to construct a nonseparable example. We now give this result, modulo various basic C^* -algebraic facts and omitting the proof of one key C^* -algebraic lemma. Whether it is relatively consistent with ZFC that no examples exist remains open.

A C^* -algebra is simple if it contains no nontrivial closed two-sided ideals. If $\mathcal{A}_1 \subseteq \mathcal{A}_2 \subseteq \dots$ is a nested sequence of simple C^* -algebras and \mathcal{A} is the completion of the union $\bigcup \mathcal{A}_n$, it is standard that \mathcal{A} is also simple. We omit the proof of the following lemma.

Lemma 16.2. *Let \mathcal{A} be a simple, separable, unital C^* -algebra and let f and g be unitarily inequivalent pure states on \mathcal{A} . Then there is a simple, separable, unital C^* -algebra \mathcal{B} that unitally contains \mathcal{A} such that f and g have unique extensions to pure states on \mathcal{B} , and these extensions are unitarily equivalent.*

(\mathcal{B} unitally contains \mathcal{A} if the unit of \mathcal{B} equals the unit of \mathcal{A} .)

Lemma 16.3. *Let (\mathcal{A}_α) , $\alpha < \aleph_1$, be a nested transfinite sequence of separable C^* -algebras and suppose \mathcal{A}_α is the completion of $\bigcup_{\beta < \alpha} \mathcal{A}_\beta$, for every limit ordinal α . Then $\mathcal{A} = \bigcup_{\alpha < \aleph_1} \mathcal{A}_\alpha$ is a C^* -algebra, and if f is a pure state on \mathcal{A} then $\{\alpha < \aleph_1 : f \text{ restricts to a pure state on } \mathcal{A}_\alpha\}$ is closed and unbounded.*

Proof. \mathcal{A} is complete because any Cauchy sequence in \mathcal{A} must lie in \mathcal{A}_α for some α , and hence must have a limit in $\mathcal{A}_\alpha \subseteq \mathcal{A}$. Using an abstract characterization of C^* -algebras, it easily follows that \mathcal{A} is a C^* -algebra.

Let f be a pure state on \mathcal{A} , let (α_n) be a sequence of countable ordinals such that $f|_{\mathcal{A}_{\alpha_n}}$ is pure for all n , and let $\alpha = \sup \alpha_n$. Then $f|_{\mathcal{A}_\alpha}$ must be pure because if $f|_{\mathcal{A}_\alpha} = (f_1 + f_2)/2$ for some states f_1 and f_2 on \mathcal{A}_α , then for all n the restrictions of f_1 and f_2 to \mathcal{A}_{α_n} must agree by purity of $f|_{\mathcal{A}_{\alpha_n}}$. Thus $f_1 = f_2$, using the fact that \mathcal{A}_α is the completion of $\bigcup \mathcal{A}_{\alpha_n}$. This shows that the set of α such that $f|_{\mathcal{A}_\alpha}$ is pure is closed. The proof that it is unbounded is essentially the same as the proof of Lemma 13.4. ■

Theorem 16.4. *Assume \diamond . Then there is a C^* -algebra generated by \aleph_1 elements that has only one irreducible representation up to unitary equivalence but is not isomorphic to the algebra of compact operators on any Hilbert space.*

Proof. By diamond, choose a sequence $\{h_\alpha : \alpha < \aleph_1\}$ of functions $h_\alpha : \alpha \rightarrow \aleph_1$ such that for any function $f : \aleph_1 \rightarrow \aleph_1$ the set $\{\alpha : f|_\alpha = h_\alpha\}$ is a stationary subset of \aleph_1 . For $\alpha < \aleph_1$ we recursively construct a nested transfinite sequence of simple separable unital C^* -algebras \mathcal{A}_α , all with the same unit, together with a pure state f_α on \mathcal{A}_α and an injective function ϕ_α from the set of states on \mathcal{A}_α into \aleph_1 . We will ensure that for any $\alpha < \beta$ the state f_α has a unique extension to a state on \mathcal{A}_β and f_β is this extension.

Begin by letting \mathcal{A}_0 be any simple, separable, infinite dimensional, unital C^* -algebra. (There are many examples of these. For instance, unitally embed $M_{2^n}(\mathbb{C}) = \mathcal{B}(\mathbb{C}^{2^n})$ in $M_{2^{n+1}}(\mathbb{C})$ by the map $A \mapsto \begin{pmatrix} A & 0 \\ 0 & A \end{pmatrix}$) and take the completion of the union. The result will be simple by a fact mentioned earlier which states that the completion of a union of a nested sequence of simple C^* -algebras is always simple.) Let f_0 be any pure state on \mathcal{A}_0 (a basic fact from convexity theory: every C^* -algebra has pure states). Since \diamond implies CH (§ 14, Exercise (b)), the set of states on \mathcal{A}_0 has cardinality at most \aleph_1 , so let ϕ_0 be any injective function from the set of states on \mathcal{A}_0 into \aleph_1 .

Having constructed \mathcal{A}_α , we define $\mathcal{A}_{\alpha+1}$ as follows. If α is a limit ordinal and there is a pure state g_α on \mathcal{A}_α that is not unitarily equivalent to f_α and such that $h_\alpha(\beta) = \phi_\beta(g_\alpha|_{\mathcal{A}_\beta})$ for all $\beta < \alpha$, then let $\mathcal{A}_{\alpha+1}$ be the C*-algebra given by Lemma 16.2 with $f = f_\alpha$ and $g = g_\alpha$. Also let $f_{\alpha+1}$ be the unique extension of f_α to $\mathcal{A}_{\alpha+1}$ and let $\phi_{\alpha+1}$ be any injective function from the set of states on $\mathcal{A}_{\alpha+1}$ into \aleph_1 . Otherwise, let $\mathcal{A}_{\alpha+1} = \mathcal{A}_\alpha$, $f_{\alpha+1} = f_\alpha$, and $\phi_{\alpha+1} = \phi_\alpha$.

At limit stages α , we let \mathcal{A}_α be the completion of $\bigcup_{\beta < \alpha} \mathcal{A}_\beta$, define f_α by the condition that its restriction to \mathcal{A}_β is f_β for all $\beta < \alpha$, and let ϕ_α be any injective function from the set of states on \mathcal{A}_α into \aleph_1 . Again, note that \mathcal{A}_α is simple because it is the completion of the union of a nested sequence of simple C*-algebras. This completes the description of the construction of the sequence (\mathcal{A}_α) .

Let $\mathcal{A} = \bigcup_{\alpha < \aleph_1} \mathcal{A}_\alpha$ and define a state f on \mathcal{A} by requiring $f|_{\mathcal{A}_\alpha} = f_\alpha$ for all α . Then \mathcal{A} is a C*-algebra by Lemma 16.3, and f is pure by the argument used in the proof of Lemma 16.3 to verify closure. We must show that every pure state on \mathcal{A} is unitarily equivalent to f . Let g be any pure state on \mathcal{A} ; we will show that $g|_{\mathcal{A}_\alpha}$ is unitarily equivalent to f_α for some α . This is sufficient because f_α has a unique state extension to \mathcal{A} (namely f), so conjugation by the unitary that makes $g|_{\mathcal{A}_\alpha}$ and f_α equivalent shows that $g|_{\mathcal{A}_\alpha}$ has a unique state extension to \mathcal{A} (namely, f), which is unitarily equivalent to f .

Define $h : \aleph_1 \rightarrow \aleph_1$ by setting $h(\alpha) = \phi_\alpha(g|_{\mathcal{A}_\alpha})$. Let S be the set of limit ordinals such that $g|_{\mathcal{A}_\alpha}$ is pure. By Lemma 16.3 together with the fact that the set of limit ordinals less than \aleph_1 is closed and unbounded, it follows that S is closed and unbounded. Thus by \diamond there exists a limit ordinal α such that $g|_{\mathcal{A}_\alpha}$ is pure and $h_\alpha = h|_{\aleph_\alpha}$, i.e.,

$$h_\alpha(\beta) = \phi_\beta(g|_{\mathcal{A}_\beta})$$

for all $\beta < \alpha$. If g_α is unitarily equivalent to f_α then we are done; otherwise the construction of $f_{\alpha+1}$ guarantees that it is unitarily equivalent to the unique extension of $g|_{\mathcal{A}_\alpha}$ to $\mathcal{A}_{\alpha+1}$, which must be $g|_{\mathcal{A}_{\alpha+1}}$. Thus we are done in this case as well.

Finally, \mathcal{A} is not isomorphic to any $\mathcal{K}(H)$ because it is both infinite dimensional and unital. ■

References

- C. Akemann and N. Weaver, Consistency of a counterexample to Naimark's problem, *Proc. Nat. Acad. Sci. USA* 101 (2004), 7522-7525.
M. A. Naimark, Rings with involutions, *Uspehi Matem. Nauk* 3 (1948), 52-145 (Russian).
———, On a problem in the theory of rings with involution, *Uspehi Matem. Nauk* 6 (1951), 160-164 (Russian).
A. Rosenberg, The number of irreducible representations of simple rings with no minimal ideals, *Amer. J. Math.* 75 (1953), 523-530.

17. Product forcing and \diamond^S

Working in ZFC*, we have developed a procedure for enlarging the given model \mathbb{M} of ZFC to another model $\mathbb{M}[G]$ using a preordered set in \mathbb{M} . But now the same procedure can be applied to $\mathbb{M}[G]$; that is, we can enlarge $\mathbb{M}[G]$ to another model of ZFC in exactly the same way, using a preordered set in $\mathbb{M}[G]$. As we will see in § 19 and subsequent sections, iterating this process increases the power of the forcing technique. In this section we will adopt the opposite point of view and see how in some circumstances the passage from \mathbb{M} to $\mathbb{M}[G]$ can be broken up into two stages. This can be useful in determining the properties of $\mathbb{M}[G]$.

Definition 17.1. Let P_1 and P_2 be two notions of forcing. The product notion of forcing $P_1 \times P_2$ is the cartesian product of P_1 and P_2 equipped with the preorder defined by $(p, q) \leq (p', q')$ if $p \leq p'$ and $q \leq q'$.

Theorem 17.2. Let P_1 and P_2 be two notions of forcing.

- (a) If G_1 is a generic filter of P_1 relative to \mathbb{M} and G_2 is a generic filter of P_2 relative to $\mathbb{M}[G_1]$, then $G = G_1 \times G_2$ is a generic filter of $P_1 \times P_2$ relative to \mathbb{M} .

(b) If G is a generic filter of $P_1 \times P_2$, then $G = G_1 \times G_2$ where G_1 is a generic filter of P_1 relative to \mathbb{M} and G_2 is a generic filter of P_2 relative to $\mathbb{M}[G_1]$.

In either case, we have $\mathbb{M}[G] = \mathbb{M}[G_1][G_2]$.

Proof. (a) It is easy to check that $G = G_1 \times G_2$ is a filter of $P_1 \times P_2$. Let $D \in \mathbb{M}$ be a dense subset of $P_1 \times P_2$; we must show that G intersects D .

Define

$$D' = \{q \in P_2 : (p, q) \in D \text{ for some } p \in G_1\}.$$

We claim that D' is dense in P_2 . To see this, let $r \in P_2$ and let

$$D'' = \{p \in P_1 : (p, q) \in D \text{ for some } q \leq r\}.$$

Since D is dense in $P_1 \times P_2$, for any $p \in P_1$ the point (p, r) lies above an element of D , and this shows that D'' is dense in P_1 . Thus G_1 intersects D'' ; fix $p \in G_1 \cap D''$ and fix $q \leq r$ such that $(p, q) \in D$. Then $q \in D'$, and this verifies the claim.

Since D' is dense in P_2 and $D' \in \mathbb{M}[G_1]$, G_2 must intersect D' . Fix $q' \in G_2 \cap D'$ and fix $p' \in G_1$ such that $(p', q') \in D$. Then $(p', q') \in G \cap D$. Thus G is generic.

(b) Let $G_1 = \{p \in P_1 : (p, 1_{P_2}) \in G\}$ and $G_2 = \{q \in P_2 : (1_{P_1}, q) \in G\}$. For any $(p, q) \in G$ we have $(p, 1_{P_2}), (1_{P_1}, q) \in G$; this shows that $G \subseteq G_1 \times G_2$. Conversely, if $p \in G_1$ and $q \in G_2$ then $(p, 1_{P_2}), (1_{P_1}, q) \in G$ and by the filter property G must contain a common extension of $(p, 1_{P_2})$ and $(1_{P_1}, q)$. But (p, q) is their greatest lower bound in $P_1 \times P_2$, so G must contain (p, q) . Thus $G_1 \times G_2 \subseteq G$. We conclude that $G = G_1 \times G_2$.

We will show that G_2 is a generic filter of P_2 relative to $\mathbb{M}[G_1]$. A similar argument shows that G_1 is a generic filter of P_1 relative to $\mathbb{M}[G_2]$, which trivially implies that it is generic relative to \mathbb{M} .

Thus, let $D \in \mathbb{M}[G_1]$ be a dense subset of P_2 . Let $\tau \in \mathbb{M}$ be a P_1 -name for D and fix $p_0 \in G_1$ such that $p_0 \Vdash \text{``}\tau \text{ is a dense subset of } \check{P}_2\text{''}$. Define

$$D' = \{(p, q) \in P_1 \times P_2 : p \leq p_0 \text{ and } p \Vdash \check{q} \in \tau\}.$$

We claim that D' is dense below $(p_0, 1_{P_2})$. To see this let $(r, s) \leq (p_0, 1_{P_2})$. Then $r \Vdash \text{``there exists } x \in \tau \text{ such that } x \leq \check{s}\text{''}$ since $r \leq p_0$. Thus some $r' \leq r$ forces $\check{s}' \in \tau$ and $\check{s}' \leq \check{s}$ for some s , so that $(r', s') \in D'$ and $(r', s') \leq (r, s)$. This proves the claim.

Since G is generic and $(p_0, 1_{P_2}) \in G$, there must exist $(p, q) \in G \cap D'$. Then $q \in G_2$, and $p \Vdash \check{q} \in \tau$ implies that $q \in D$. This shows that G_2 is generic relative to $\mathbb{M}[G_1]$.

For the final statement of the theorem, observe that G_1, G_2 , and G all belong to both $\mathbb{M}[G]$ and $\mathbb{M}[G_1][G_2]$. The fact that $\mathbb{M}[G] \subseteq \mathbb{M}[G_1][G_2]$ is then a straightforward induction on name rank, as is the reverse containment. ■

We will now illustrate the way product decompositions can be employed by using them to prove a slightly stronger version of the diamond principle.

Definition 17.3. For any stationary set $S \subseteq \aleph_1$, $\diamond(S)$ is the assertion that there is a sequence $\{h_\alpha : \alpha \in S\}$ of functions $h_\alpha : \alpha \rightarrow \aleph_1$ such that for any function $f : \aleph_1 \rightarrow \aleph_1$ the set $\{\alpha \in S : f|_\alpha = h_\alpha\}$ is stationary. \diamond^S is the assertion that $\diamond(S)$ holds for every stationary $S \subseteq \aleph_1$.

Thus, $\diamond(S)$ says that paths down the standard \aleph_1 - \aleph_1 -tree can be blocked just by removing one vertex from each level in S , and \diamond^S says this can be done for every stationary set S .

We will force \diamond^S by adding not just one diamond sequence as in Theorem 14.2, but more than 2^{\aleph_0} diamond sequences. The notion of forcing used in Theorem 14.2 actually verifies $\diamond(S)$ for every stationary S in \mathbb{M} , so the idea is that if we add many diamond sequences at once then for every stationary S in $\mathbb{M}[G]$, some of

the diamond sequences we create are added “independently” of S and hence will verify $\diamond(S)$ for that S . The technical tool used to show this is a product decomposition. (This argument follows a suggestion of Ilijas Farah.)

First we show that adding a single generic diamond sequence already verifies $\diamond(S)$ for every S in the ground model.

Lemma 17.4. *Let $S \in \mathbb{M}$ be a subset of \aleph_1 such that $\mathbb{M} \models$ “ S is stationary” and let P be the notion of forcing used in the proof of Theorem 14.2. Then $\diamond(S)$ is true in $\mathbb{M}[G]$, for any generic filter G .*

Proof. Fix a generic filter G and suppose $\diamond(S)$ fails in $\mathbb{M}[G]$. Let (h_α) be the sequence of functions $h_\alpha : \alpha \rightarrow \aleph_1^{\mathbb{M}}$ which constitute the union of G . Then we can find a function f in $\mathbb{M}[G]$ from $\aleph_1^{\mathbb{M}}$ to itself and a subset $C \subseteq \aleph_1^{\mathbb{M}}$ in $\mathbb{M}[G]$ such that $\mathbb{M} \models$ “ C is closed and unbounded, and $f|_C \neq h_\alpha$ for any $\alpha \in C \cap S$ ”. Let τ be a P -name for f and let σ be a P -name for C , and find $p \in G$ such that p forces “ τ is a function from \aleph_1 to itself, σ is a closed and unbounded subset of \aleph_1 , and τ restricted to α is not in the union of Γ for any $\alpha \in \sigma \cap \check{S}$ ”.

Working in \mathbb{M} , carry out the construction given in the proof of Theorem 14.2 to obtain $q_0 < p$ of length α_0 such that q_0 forces “ $\check{\alpha}_0 \in \sigma$ and τ restricted to $\check{\alpha}_0$ is in the union of Γ ”. Then carry out the construction again to get $q_1 < q_0$ of length α_1 with the same property, and iterate. We get a sequence $\{\alpha_\nu : \nu \in \aleph_1^{\mathbb{M}}\}$ which enumerates a closed unbounded subset D of \aleph_1 , together with elements $q_\nu \in P$ which force that $\check{\alpha}_\nu$ is in σ and τ restricted to $\check{\alpha}_\nu$ is in the union of Γ . Since S is stationary, $S \cap D \neq \emptyset$, so let $\alpha_{\nu^*} \in S \cap D$. We now reach a contradiction, because q_{ν^*} forces that Γ captures τ at the level $\check{\alpha}_{\nu^*}$, but since it lies below p it also forces that Γ does not capture τ at any level in $\sigma \cap \check{S}$. We conclude that $\diamond(S)$ must hold in $\mathbb{M}[G]$. ■

Next we need a version of the Δ -systems lemma (Lemma 10.5) for countable sets.

Lemma 17.5. *Say $2^{\aleph_0} = \aleph_\alpha$ and let $\kappa = \aleph_{\alpha+1}$. Let A be a family of κ distinct countable subsets of a set of cardinality κ . Then there is a set r and a subfamily $B \subseteq A$ of cardinality κ such that $a \cap b = r$ for any distinct $a, b \in B$.*

Proof. Without loss of generality, suppose A is a family of subsets of the cardinal κ . We claim that there exists an ordinal $\gamma < \kappa$ such that for every $\delta > \gamma$, there is some set $a \in A$ which contains an ordinal larger than δ but no ordinals between γ and δ .

Suppose the claim fails. Then for every $\gamma < \kappa$ there exists $\delta > \gamma$ such that every set in A either contains no ordinal larger than γ or contains an ordinal between γ and δ . Create an increasing sequence γ_β , $\beta < \aleph_1$, by setting $\gamma_0 = 0$ and using the failure of the claim to find $\gamma_{\beta+1} = \delta$ for every $\gamma_\beta = \gamma$.

For every $\gamma < \kappa$ we have $\text{card}(\gamma) \leq 2^{\aleph_0}$, so there are at most 2^{\aleph_0} distinct countable subsets of γ (§ 3, Exercise (e)). So for each $\beta < \aleph_1$, there are at most 2^{\aleph_0} sets in A contained in γ_β , and hence there are at most $\aleph_1 \cdot 2^{\aleph_0} = 2^{\aleph_0}$ sets in A that are contained in any γ_β . Since A has cardinality $\kappa > 2^{\aleph_0}$, there must be some $a \in A$ that is not contained in any γ_β . But then by the construction of $\gamma_{\beta+1}$ we have that a contains an ordinal between γ_β and $\gamma_{\beta+1}$ for all $\beta < \aleph_1$, which is impossible since every element of A is countable. This contradiction establishes the claim.

Now fix $\gamma < \kappa$ verifying the claim. Choose a sequence of sets $a_\beta \in A$, $\beta < \kappa$, by picking a_0 arbitrarily and for general β using the lemma to find a set $a_\beta \in A$ which contains no ordinal between γ and $(\sup \bigcup_{\beta' < \beta} a_{\beta'}) + 1$. Then $\{a_\beta : \beta < \kappa\}$ is a subfamily of A of κ distinct countable sets, the intersection of any two of which is a countable subset of γ . But as we noted earlier, there are at most $2^{\aleph_0} < \kappa$ countable subsets of γ , so for some countable $r \subseteq \gamma$ we have $a \cap b = r$ for all distinct a and b in some subfamily B of $\{a_\beta : \beta < \kappa\}$ of cardinality κ . ■

Lemma 17.6. *Let P be any notion of forcing, let G be a generic filter of P , let A be any set in \mathbb{M} , and let τ be a P -name such that $\mathbb{M}[G] \models$ “ $\text{val}_G(\tau) \subseteq A$ ”. Then there is a P -name σ whose elements are all of the form $\langle \check{x}, p \rangle$ for $x \in A$ and $p \in P$, such that $\text{val}_G(\tau) = \text{val}_G(\sigma)$. If $\mathbb{M} \models$ “ $\text{card}(A) \leq 2^{\aleph_0}$ and every*

down-antichain in P has cardinality at most 2^{\aleph_0} ” then we can choose σ so that $\mathbb{M} \models \text{“}\sigma \text{ contains at most } 2^{\aleph_0} \text{ elements”}$.

Proof. Let σ be the set of all pairs $\langle \check{x}, p \rangle$ such that $x \in A$, $p \in P$, and $p \Vdash \check{x} \in \tau$. It is an exercise to verify that $\text{val}(\tau) = \text{val}(\sigma)$.

To prove the second assertion, working in \mathbb{M} , for each $x \in A$ let S_x be the set of $p \in P$ which force $\check{x} \in \tau$, and let T_x be a maximal down-antichain in S_x . Then let σ' be the set of all pairs $\langle \check{x}, p \rangle$ such that $x \in A$ and $p \in T_x$. If $\mathbb{M} \models \text{“card}(A) \text{ and card}(T_x) \text{ are both at most } 2^{\aleph_0}\text{”}$ for all $x \in A$, then $\mathbb{M} \models \text{“}\sigma' \text{ contains at most } 2^{\aleph_0} \text{ elements”}$. The fact that $\text{val}(\tau) = \text{val}(\sigma')$ is another exercise. \blacksquare

The P -name σ' in Lemma 17.6 is called a nice name for a subset of A .

Theorem 17.7. *Say $\mathbb{M} \models \text{“}2^{\aleph_0} = \aleph_\alpha \text{ and } \kappa = \aleph_{\alpha+1}\text{”}$. Let P be the notion of forcing defined in \mathbb{M} as all families of sequences $\{f_\gamma^\delta : \delta \in A \text{ and } \gamma < \alpha\}$ such that A is a countable subset of κ , $\alpha < \aleph_1$, and each f_γ^δ is a function from γ into \aleph_1 . Then \diamond^S is true in $\mathbb{M}[G]$, for any generic filter G of P .*

Proof. Fix $S \in \mathbb{M}[G]$ such that $\mathbb{M}[G] \models \text{“}S \text{ is a stationary subset of } \aleph_1\text{”}$ and let τ be a P -name for S . It follows from Lemma 17.5 that $\mathbb{M} \models \text{“}P \text{ contains no down-antichains of cardinality greater than } 2^{\aleph_0}\text{”}$ (exercise), so we can use Lemma 17.6 with $A = \aleph_1$ to find a nice name σ for S such that $\mathbb{M} \models \text{“}\sigma \text{ contains at most } 2^{\aleph_0} \text{ elements”}$.

Working in \mathbb{M} , for every pair $\langle \check{x}, p \rangle$ in σ let A_p be the countable subset of κ such that p is a family of sequences f_γ^δ with $\delta \in A_p$. Since σ contains at most 2^{\aleph_0} pairs, the union of all the A_p has cardinality at most 2^{\aleph_0} . So there must exist an ordinal $\delta_0 < \kappa$ that does not belong to any A_p .

We have $P \cong P_1 \times P_2$ where P_1 consists of all elements of P for which $\delta_0 \notin A$, and P_2 consists of all elements of P for which $A = \{\delta_0\}$. Let $G \cong G_1 \times G_2$ be the corresponding decomposition of G . Then $\text{val}_{G_1}(\sigma) = \text{val}_G(\sigma)$ since any pair $\langle \check{x}, p \rangle$ in σ satisfies $p \in P_1$, and hence $S \in \mathbb{M}[G_1]$. Since $\mathbb{M}[G] \models \text{“}S \text{ is stationary”}$ it follows that $\mathbb{M}[G_1] \models \text{“}S \text{ is stationary”}$ (exercise). Then $\mathbb{M}[G] = \mathbb{M}[G_1][G_2]$ satisfies $\diamond(S)$ by Lemma 17.4, and as S was arbitrary we conclude that \diamond^S is true in $\mathbb{M}[G]$. \blacksquare

Exercises

- (a) In the proof of Lemma 17.6, verify that $\text{val}(\tau) = \text{val}(\sigma)$ and $\text{val}(\tau) = \text{val}(\sigma')$.
- (b) In the proof of Theorem 17.7, verify that $\mathbb{M} \models \text{“}P \text{ contains no down-antichains of cardinality greater than } 2^{\aleph_0}\text{”}$.
- (c) In the proof of Theorem 17.7, show that $\mathbb{M}[G] \models \text{“}S \text{ is stationary”}$ implies that $\mathbb{M}[G_1] \models \text{“}S \text{ is stationary”}$. (Hint: show that if $\mathbb{M}[G_1] \models \text{“}C \text{ is closed and unbounded”}$ then $\mathbb{M}[G] \models \text{“}C \text{ is closed and unbounded”}$.)

18. Application: The Whitehead problem, I

We work in ZFC. In this section all groups are abelian.

Definition 18.1. Let A be an abelian group.

- (a) A is free if there is a set $\{g_i : i \in I\}$ of elements of A (a basis for A) such that every $g \in A$ is uniquely expressible in the form

$$g = n_1 g_{i_1} + \cdots + n_k g_{i_k}$$

with $i_1, \dots, i_k \in I$ and $n_1, \dots, n_k \in \mathbb{Z}$.

- (b) An extension of A by \mathbb{Z} is an abelian group B together with a surjective homomorphism $\pi : B \rightarrow A$ with kernel isomorphic to \mathbb{Z} , i.e., it is a short exact sequence

$$0 \rightarrow \mathbb{Z} \rightarrow B \rightarrow A \rightarrow 0.$$

It is trivial if $B \cong A \oplus \mathbf{Z}$ and π is projection onto the first summand.

Proposition 18.2. *Let A be a free abelian group. Then every extension of A by \mathbf{Z} is trivial.*

Proof. Let $\pi : B \rightarrow A$ be a surjective homomorphism with kernel \mathbf{Z} . Fix a basis $\{g_i : i \in I\}$ for A , and for each $i \in I$ choose $h_i \in \pi^{-1}(g_i)$. Let A' be the subgroup of B generated by the h_i . Then it is straightforward to check that $A' \cong A$, $A' \cap \ker \pi = \{0\}$, and $A' + \ker \pi = B$. Thus $B \cong A \oplus \mathbf{Z}$ and π is projection onto the first summand. \blacksquare

The idea of this proposition can be expressed more abstractly. An extension $\pi : B \rightarrow A$ of A by \mathbf{Z} is trivial if and only if there is a homomorphism $\rho : A \rightarrow B$ such that $\pi \circ \rho = \text{id}_A$. (If $B \cong A \oplus \mathbf{Z}$ and π is projection onto the first factor, the definition of ρ is obvious; if there is such a map ρ , then $B = \rho(A) + \ker \pi \cong A \oplus \mathbf{Z}$.) We say that ρ splits π .

According to Proposition 18.2, every free abelian group has only trivial extensions by \mathbf{Z} . The Whitehead problem asks: if A has only trivial extensions by \mathbf{Z} , is it free?

The argument of Proposition 18.2 actually shows that any extension of a free abelian group by any abelian group is trivial, and the converse of this statement is true in ZFC: if all extensions of A by any abelian group are trivial, then A is free. But when A is countable, having only trivial extensions by \mathbf{Z} is enough:

Theorem 18.3. *Let A be a countable abelian group and suppose all extensions of A by \mathbf{Z} are trivial. Then A is free.*

We omit the proof of Theorem 18.3 (see Exercise (c)). Using this result, we will show \diamond^S implies the same result when A has cardinality \aleph_1 , a theorem due to Saharon Shelah. Shelah also proved that versions of \diamond^S for arbitrary cardinals give the result for groups of arbitrary cardinality. Thus, it is relatively consistent with ZFC that any abelian group with only trivial extensions by \mathbf{Z} is free.

The next lemma can be proven using techniques from homological algebra, but we give an elementary proof. This proof uses an idea of Mohan Kumar.

Lemma 18.4. *Let A be an abelian group, let A_0 be a subgroup of A , and let $\pi_0 : B_0 \rightarrow A_0$ be an extension of A_0 by \mathbf{Z} . Then π_0 embeds in an extension $\pi : B \rightarrow A$ of A by \mathbf{Z} .*

Proof. Using Zorn's lemma, it is enough to consider the case that A is generated by A_0 and one additional element a . Suppose this is the case. If the cyclic subgroup of A generated by a does not intersect A_0 then $A \cong A_0 \oplus \mathbf{Z}$ and we can define $B = B_0 \oplus \mathbf{Z}$ and $\pi(x \oplus m) = \pi_0(x) \oplus m$.

Otherwise, let n be the least positive integer such that $na \in A_0$, and write $b = na$. Then fix $b' \in B_0$ such that $\pi_0(b') = b$ and let $B = B_0 \oplus \mathbf{Z}/I$ where I is the cyclic subgroup of $B_0 \oplus \mathbf{Z}$ generated by $b' \oplus n$. Define $\pi : B \rightarrow A$ by $\pi((x \oplus m) + I) = \pi_0(x) - ma$. This is well-defined since $\pi_0(b') - na = b - na = 0$. We can embed B_0 in B by the map $x \mapsto (x \oplus 0) + I$, and the restriction of π to B_0 then agrees with π_0 . The kernel of π consists of all $(x \oplus m) + I \in B$ such that $\pi_0(x) = ma$. Since $\pi_0(x) \in A_0$ we must have $m = nk$ for some integer k , and then $\pi_0(x) = ma = nka = kb$ implies that $x = kb' + z$ for some $z \in \ker \pi_0$. That is,

$$(x \oplus m) + I = (kb' + z \oplus nk) + I = (z \oplus 0) + I$$

since $kb' \oplus nk = k(b' \oplus n) \in I$, so that $\ker \pi = \ker \pi_0 = \mathbf{Z}$. So $\pi : B \rightarrow A$ is the desired extension of A by \mathbf{Z} . \blacksquare

It immediately follows from Theorem 18.3 and Lemma 18.4 that if A has only trivial extensions by \mathbf{Z} then all countable subgroups of A are free.

Lemma 18.5. *Let A be an abelian group, A_0 a subgroup of A , and $\rho_0 : A_0 \rightarrow A_0 \oplus \mathbf{Z}$ the standard splitting of the trivial extension of A_0 by \mathbf{Z} . Suppose that the quotient group A/A_0 has a nontrivial extension by \mathbf{Z} . Then the trivial extension of A_0 by \mathbf{Z} can be embedded in an extension of A by \mathbf{Z} no splitting of which restricts to ρ_0 on A_0 .*

Proof. Let $A' = A/A_0$ and let $\pi' : B' \rightarrow A'$ be a nontrivial extension of A' by \mathbf{Z} . Let B be the subgroup of $A \oplus B'$ consisting of all elements of the form $x \oplus y$ such that $x + A_0 = \pi'(y)$, and define $\pi : B \rightarrow A$ by $\pi(x \oplus y) = x$. Then π is surjective because π' is surjective: for every $x \in A$ there exists $y \in B'$ such that $\pi'(y) = x + A_0$. The kernel of π consists of all $x \oplus y$ such that $x = 0$ and $\pi'(y) = A_0$, i.e., it consists of all elements of the form $0 \oplus y$ for $y \in \ker \pi'$. So $\ker \pi \cong \ker \pi' = \mathbf{Z}$. Thus $\pi : B \rightarrow A$ is an extension of A by \mathbf{Z} .

B contains $A_0 \oplus \mathbf{Z}$, and restricting π to this subgroup yields the trivial extension of A_0 by \mathbf{Z} . Suppose the standard splitting $\rho_0 : A_0 \rightarrow A_0 \oplus \mathbf{Z}$ defined by $\rho_0(x) = x \oplus 0$ extended to a splitting $\rho : A \rightarrow B$ of π . Then we could define a splitting $\rho' : A' \rightarrow B'$ of π' by setting $\rho'(x + A_0) = y$ where $\rho(x) = x \oplus y$. This would be well-defined because for $x \in A_0$ we would have $\rho(x) = \rho_0(x) = x \oplus 0$, and it would be a splitting because $x \oplus y \in B$ implies $\pi'(y) = x + A_0$. This is impossible because π' is a nontrivial extension of A' , so we conclude that ρ_0 does not extend to a splitting of π . ■

The next lemma contains the substance of the proof.

Lemma 18.6. *Assume \diamond^S and let A be a group of cardinality \aleph_1 all of whose extensions by \mathbf{Z} are trivial. Suppose $A = \bigcup_{1 \leq \alpha < \aleph_1} A_\alpha$ where the A_α are strictly nested countably infinite subgroups of A such that $A_\alpha = \bigcup_{\beta < \alpha} A_\beta$ when α is a limit ordinal. Then $A_{\alpha+1}/A_\alpha$ is free for all α in some closed unbounded set $C \subseteq \aleph_1$.*

Proof. Construct a tree as follows. The vertices at level α are all functions $f : A_\alpha \rightarrow \mathbf{N} \times \alpha$ such that $f(A_\beta) \subseteq \mathbf{N} \times \beta$ for all $\beta < \alpha$, with g lying below f if g is an extension of f . Suppose the lemma fails; then the set $S \subseteq \aleph_1$ of values of α such that $A_{\alpha+1}/A_\alpha$ is not free is stationary. Use \diamond^S to choose a sequence of vertices $h_\alpha, \alpha \in S$.

We now recursively construct a nested sequence of countable abelian groups $C_\alpha, \alpha < \aleph_1$, such that the underlying set of C_α is the set $\mathbf{N} \times \alpha$, together with surjections $\pi_\alpha : C_\alpha \rightarrow A_\alpha$ satisfying $\ker \pi_\alpha \cong \mathbf{Z}$ and $\pi_\alpha|_{C_\beta} = \pi_\beta$ for all $\beta < \alpha$.

Begin the construction by letting $\pi_1 : C_1 \rightarrow A_1$ be the trivial extension of A_1 by \mathbf{Z} and identifying the underlying set of C_1 with \mathbf{N} in any way. At limit levels let $C_\alpha = \bigcup_{\beta < \alpha} C_\beta$ and let π_α be the direct limit of the $\pi_\beta, \beta < \alpha$. At successor levels, we have two cases. If $\alpha \in S$ and h_α splits π_α , use the fact that $A_{\alpha+1}/A_\alpha$ is not free (since $\alpha \in S$) plus Theorem 18.3 to find a nontrivial extension of $A_{\alpha+1}/A_\alpha$ by \mathbf{Z} . Then apply Lemma 18.5 to get an extension $\pi_{\alpha+1} : C_{\alpha+1} \rightarrow A_{\alpha+1}$ by \mathbf{Z} which restricts to π_α on C_α but such that h_α does not extend to a splitting of $\pi_{\alpha+1}$. We may arrange that the underlying set of $C_{\alpha+1}$ is $\mathbf{N} \times (\alpha + 1)$.

The other case is that $\alpha \notin S$ or h_α fails to split π_α . Here we use Lemma 18.4 to embed $\pi_\alpha : C_\alpha \rightarrow A_\alpha$ in any extension $\pi_{\alpha+1} : C_{\alpha+1} \rightarrow A_{\alpha+1}$ of $A_{\alpha+1}$ by \mathbf{Z} . Again, we may arrange that the underlying set of $C_{\alpha+1}$ is $\mathbf{N} \times (\alpha + 1)$. This completes the description of the construction.

Let $C = \bigcup_{1 \leq \alpha < \aleph_1} C_\alpha$ and let $\pi : C \rightarrow A$ be the union of the maps π_α . This is an extension of A by \mathbf{Z} , and since A has only trivial extensions by \mathbf{Z} , there is a splitting $\rho : A \rightarrow C$ of π . We have $\rho(A_\alpha) \subseteq C_\alpha$ for all α (exercise). By \diamond^S there therefore exists $\alpha \in S$ such that $\rho|_{A_\alpha} = h_\alpha$. This puts us in the first case at that stage of the construction, so that h_α does not extend to a splitting of $\pi_{\alpha+1}$. But $\rho|_{A_{\alpha+1}}$ is just such a splitting, which is a contradiction. Thus $A_{\alpha+1}/A_\alpha$ could not fail to be free on a stationary set. ■

Lemma 18.7. *Assume \diamond^S and let A be a group of cardinality \aleph_1 all of whose extensions by \mathbf{Z} are trivial. Let A_1 be a countably infinite subgroup of A . Then there exists a countable subgroup B_1 of A containing A_1 such that for any countable subgroup B of A containing B_1 the quotient B/B_1 is free.*

Proof. Suppose the lemma fails. Then

(†) for any countable subgroup B_1 containing A_1 there exists a countable subgroup B containing B_1 such that B/B_1 is not free.

Enumerate the elements of A as $\{x_\alpha : \alpha < \aleph_1\}$. Now recursively construct a strictly nested sequence of countable subgroups A_α of A such that (1) $A_\alpha = \bigcup_{\beta < \alpha} A_\beta$ when α is a limit ordinal; (2) using (†), $A_{\alpha+1}/A_\alpha$

is not free for any limit ordinal α ; and (3) $x_\alpha \in A_{\alpha+2}$ for all $\alpha < \aleph_1$. By (3) we get $A = \bigcup_{1 \leq \alpha < \aleph_1} A_\alpha$, and then (2) contradicts Lemma 18.6. \blacksquare

Theorem 18.8. *Assume \diamond^S and let A be a group of cardinality \aleph_1 all of whose extensions by \mathbf{Z} are trivial. Then A is free.*

Proof. By Lemma 18.7 we can decompose A as a nested transfinite sequence (A_α) , $1 \leq \alpha < \aleph_1$, in such a way that $A_{\alpha+1}$ is always chosen to have the property that $B/A_{\alpha+1}$ is free, and thus $B \cong A_{\alpha+1} \oplus B/A_{\alpha+1}$, for any countable subgroup B of A that contains $A_{\alpha+1}$. We cannot immediately guarantee this property at limit stages, but by Lemma 18.6 we can find a closed unbounded subset $C \subseteq \aleph_1$ such that $A_{\alpha+1}/A_\alpha$ is free for all $\alpha \in C$, so that $A_{\alpha+1} \cong A_\alpha \oplus A_{\alpha+1}/A_\alpha$. It follows that for any $\beta \geq \alpha + 1$ we have

$$A_\beta \cong A_{\alpha+1} \oplus A_\beta/A_{\alpha+1} \cong A_\alpha \oplus A_{\alpha+1}/A_\alpha \oplus A_\beta/A_{\alpha+1}$$

so that $A_\beta/A_\alpha \cong A_{\alpha+1}/A_\alpha \oplus A_\beta/A_{\alpha+1}$ is free. Now C is order-isomorphic to \aleph_1 , so by discarding indices not in C and relabelling we may assume that $C = \aleph_1$. We then have that A_β/A_α is free whenever $\beta > \alpha$.

We can now recursively construct a sequence of subgroups B_α of A such that $A_{\alpha+1} = A_\alpha \oplus B_\alpha$ for all α . Letting $A_0 = \{0\}$, we then have $A = \bigoplus_{\alpha < \aleph_1} B_\alpha$ where each B_α is free, and this implies that A is free. \blacksquare

Exercises

- (a) In the proof of Lemma 18.6, verify that $\rho(A_\alpha) \subseteq C_\alpha$ for all α .
- (b) Let A be a finitely generated abelian group and suppose all extensions of A by \mathbf{Z} are trivial. Prove that A is free. (Use the structure theorem for finitely generated abelian groups.)
- (c) Assume the countable version of Lemma 18.7: if A is a countable abelian group with only trivial extensions by \mathbf{Z} , then any finitely generated subgroup A_1 of A is contained in a finitely generated subgroup B_1 such that B/B_1 is free for any finitely generated subgroup B containing B_1 . Then prove Theorem 18.3. (Hint: adapt the proof of Theorem 18.8, using Exercise (b) in place of Theorem 18.3.)

References

- P. Eklof, *Set Theoretic Methods in Homological Algebra and Abelian Groups*, 1980.
- S. Shelah, Infinite abelian groups, Whitehead problem and some constructions, *Israel J. Math.* 18 (1974), 243-256.

19. Two-stage iterated forcing

A two-stage iterated forcing construction involves first choosing a notion of forcing P in \mathbf{M} and constructing $\mathbf{M}[G]$ where G is a generic filter of P , then choosing a notion of forcing Q in $\mathbf{M}[G]$ and constructing $\mathbf{M}[G][H]$ where H is a generic filter of Q . Using a name for Q in \mathbf{M} , it is always possible to compress this procedure into a single step.

Definition 19.1. Let P be a notion of forcing in \mathbf{M} and let π, \leq_π , and 1_π be P -names such that

$$1_P \Vdash \text{“}\leq_\pi \text{ is a preorder of } \pi \text{ with greatest element } 1_\pi\text{”}.$$

Then $P * \pi$ is the notion of forcing whose underlying set is

$$\{\langle p, \tau \rangle : p \in P, \tau \in \text{dom}(\pi), \text{ and } p \Vdash \tau \in \pi\}$$

and with $\langle p, \tau \rangle \leq \langle q, \sigma \rangle$ if $p \leq q$ and $p \Vdash \tau \leq_\pi \sigma$. The greatest element of $P * \pi$ is $\langle 1_P, 1_\pi \rangle$.

In general $P * \pi$ is not a partially ordered set, even if P is partially ordered, because it is possible that p forces $\tau = \sigma$ for distinct P -names τ and σ , and this would imply $\langle p, \tau \rangle \leq \langle p, \sigma \rangle$ and $\langle p, \sigma \rangle \leq \langle p, \tau \rangle$.

Proposition 19.2. *Let P and π be as in Definition 19.1 and let K be a filter of $P * \pi$ which is generic relative to M . Let*

$$G = \{p \in P : \langle p, 1_\pi \rangle \in K\}$$

and

$$H = \{\text{val}_G(\tau) : \langle 1_P, \tau \rangle \in K\}.$$

Then G is a filter of P which is generic relative to M , H is a filter of $\text{val}_G(\pi)$ which is generic relative to $M[G]$, and $M[K] = M[G][H]$.

Proof. The fact that G is a filter of P which is generic relative to M is an exercise. To see that H is a filter of $\text{val}_G(\pi)$, let $\langle 1_P, \tau \rangle \in K$ and suppose $\text{val}_G(\tau) \leq \text{val}_G(\sigma)$. Then some $p \in G$ forces $\tau \leq \sigma$, and since $p \in G$ we have $\langle p, 1_\pi \rangle \in K$. Since K is a filter there must exist $\langle p', \tau' \rangle \in K$ less than both $\langle 1_P, \tau \rangle$ and $\langle p, 1_\pi \rangle$; then p' forces $\tau' \leq \tau$ and (since it is less than p) $\tau \leq \sigma$, so p' forces $\tau' \leq \sigma$ and hence $\langle p', \tau' \rangle \leq \langle 1_P, \sigma \rangle$. We conclude that $\langle 1_P, \sigma \rangle \in K$ and hence $\text{val}_G(\sigma) \in H$. This shows that H is upwards closed. It is directed downwards because if $\langle 1_P, \tau \rangle \in K$ and $\langle 1_P, \sigma \rangle \in K$ then there exists $\langle p, \rho \rangle \in K$ such that p forces $\rho \leq \tau$ and $\rho \leq \sigma$. Then $\langle 1_P, \rho \rangle \in K$ since K is upwards closed, so that $\text{val}_G(\rho) \in H$, and $\text{val}_G(\rho)$ is less than both $\text{val}_G(\tau)$ and $\text{val}_G(\sigma)$ because $p \in G$. So H is directed downwards.

The proof that H intersects every dense subset of $\text{val}_G(\pi)$ in $M[G]$ is essentially the same as the proof that G_2 is generic relative to $M[G_1]$ in Theorem 17.2 (b). The proof that $M[K] = M[G][H]$ is essentially the same as the proof that $M[G] = M[G_1][G_2]$ in Theorem 17.2. \blacksquare

We will need the following fact about two-stage iterated forcing.

Proposition 19.3. *Let P and π be as in Definition 19.1. Suppose $M \models$ “ P is c.c.c.” and $1_P \Vdash$ “ π is c.c.c.” Then $M \models$ “ $P * \pi$ is c.c.c.”*

Proof. Let γ be an ordinal in M and suppose $\{\langle p_\alpha, \tau_\alpha \rangle : \alpha < \gamma\}$ is an antichain in $P * \pi$ in M . Let G be a filter of P which is generic relative to M and let σ be the P -name consisting of all pairs $\langle \check{\alpha}, p_\alpha \rangle$. We claim that the elements $\text{val}_G(\tau_\alpha)$ for $\alpha \in \text{val}_G(\sigma)$ constitute an antichain in $\text{val}_G(\pi)$. Indeed, if $\text{val}_G(\tau_\alpha)$ and $\text{val}_G(\tau_\beta)$ had a common extension $\text{val}_G(\rho)$ for some $\alpha, \beta \in \text{val}_G(\sigma)$, then we could find $p \in G$ which forces $\rho \leq \tau_\alpha$ and $\rho \leq \tau_\beta$. Since $\alpha, \beta \in \text{val}_G(\sigma)$ we have $p_\alpha, p_\beta \in G$, so we may assume $p \leq p_\alpha$ and $p \leq p_\beta$. But then $\langle p, \rho \rangle$ must be a common extension of $\langle p_\alpha, \tau_\alpha \rangle$ and $\langle p_\beta, \tau_\beta \rangle$ in $P * \pi$, a contradiction. This proves the claim.

Since 1_P forces “ π is c.c.c.” it must therefore force “ σ is countable”. Now, working in M , for each $\alpha < \gamma$ choose $q_\alpha \in P$ which forces “ $\check{\alpha} = \text{sup}(\sigma)$ ”, if any such q_α exists. Then the set of all q_α is an antichain in P in M , so since $M \models$ “ P is c.c.c.” we must have that $M \models$ “ β is countable” where β is the supremum of all α which are forced by some element of P to equal $\text{sup}(\sigma)$. Since $p_\alpha \Vdash$ “ $\check{\alpha} \in \sigma$ ”, it follows that every $\alpha < \gamma$ is less than β , i.e., $\gamma \leq \beta$. Thus $M \models$ “ γ is countable”, so we conclude that $M \models$ “ $P * \pi$ is c.c.c.” \blacksquare

Exercises

- In Proposition 19.2, prove that G is a filter of P which is generic relative to M . Also check that H is generic relative to $M[G]$.
- Let $\langle p, \tau \rangle \in P * \pi$ and $q \in P$. Prove that p is incompatible with q in P if and only if $\langle p, \tau \rangle$ is incompatible with $\langle q, 1_\pi \rangle$ in $P * \pi$.

20. Finite support iterations

We now generalize two-stage iterated forcing to sequential forcing constructions of arbitrary length. Typically this technique is used in such a way that each stage of the forcing construction destroys one particular counterexample of some sort, and the entire sequence of forcing constructions results in a model in which there are no counterexamples of the kind in question. In order for this to work we need to be able to destroy any single potential counterexample by forcing, and we also need to be able to arrange the sequential

construction so that all counterexamples are eventually handled, and no new counterexamples appear in the final model.

In iterated forcing constructions, successor stages always follow the two-stage construction, but there are different ways of handling limit stages. Here we present the most standard construction using “finite supports”.

Definition 20.1. Let α be an ordinal in M . An α -stage finite support iterated forcing construction consists of (1) a sequence $\{P_\gamma : \gamma \leq \alpha\}$ in M such that each P_γ is a notion of forcing with preorder \leq_γ and greatest element 1_γ , and (2) a sequence $\{\pi_\gamma : \gamma < \alpha\}$ in M such that each π_γ is a P_γ -name, together with P_γ -names \leq'_γ and $1'_\gamma$ such that

$$1_\gamma \Vdash \text{“}\leq'_\gamma \text{ is a preorder of } \pi_\gamma \text{ with greatest element } 1'_\gamma\text{”}.$$

For every γ , every element of P_γ must be a sequence $\vec{\rho} = \{\rho_\mu : \mu < \gamma\}$ with $\rho_\mu \in \text{dom}(\pi_\mu)$. (Thus P_0 has only one element, the empty sequence.) At successor stages, we require $P_{\gamma+1} = P_\gamma * \pi_\gamma$, so that its elements are just those sequences $\vec{\rho}$ of length $\gamma + 1$ such that the truncated sequence $\{\rho_\mu : \mu < \gamma\}$ belongs to P_γ and forces $\rho_\gamma \in \pi_\gamma$, and the order relation is likewise as described in Definition 19.1. At limit stages we require that P_γ consist of precisely those sequences $\vec{\rho}$ such that $\rho_\mu = 1'_\mu$ for all but finitely many $\mu < \gamma$, and $\vec{\rho}|_\mu \in P_\mu$ for all $\mu < \gamma$. It is ordered by setting $\vec{\rho} \leq_\gamma \vec{\rho}'$ if $\vec{\rho}|_\mu \leq_\mu \vec{\rho}'|_\mu$ for all $\mu < \gamma$.

The next result is a partial analog of Proposition 19.2.

Proposition 20.2. *Given an α -stage finite support iterated forcing construction as in Definition 20.1, let G be a filter of P_α which is generic relative to M . Then for every $\gamma < \alpha$, the set G_γ of elements of P_γ whose extension by $\{1_\mu : \mu \geq \gamma\}$ belongs to G is a filter of P_γ which is generic relative to M . Also, for every $\gamma < \alpha$ the set H_γ of all $\text{val}(\rho)$ such that $\rho \in \text{dom}(\pi_\gamma)$ and the sequence $\{1_\mu : \mu < \gamma\}$ followed by ρ belongs to $G_{\gamma+1}$ is a filter of $\text{val}_{G_\gamma}(\pi_\gamma)$ which belongs to $M[G_{\gamma+1}]$ and is generic relative to $M[G_\gamma]$.*

Proof. Exercise. ■

We also need the following two facts. We state the second one only in the special case that we need, $\alpha = \aleph_2$, but it can be generalized to the case that α is any regular cardinal. (A cardinal κ is regular if a set of size κ cannot be written as the union of fewer than κ sets each of size less than κ . Every infinite successor cardinal, i.e., every $\aleph_{\alpha+1}$, is regular; \aleph_ω is the smallest infinite cardinal that is not regular.)

Proposition 20.3. *Suppose $1_\gamma \Vdash \text{“}\pi_\gamma \text{ is c.c.c.”}$ for all $\gamma < \alpha$. Then $M \models \text{“}P_\alpha \text{ is c.c.c.”}$*

Proof. Working in M , we prove that P_γ is c.c.c. by induction on γ . At successor stages this follows from Proposition 19.3. Now suppose γ is a limit ordinal and let A be an uncountable subset of P_γ . For $\vec{\rho} \in A$, say that $\{\gamma : \rho_\gamma \neq 1'_\gamma\}$ is the support of $\vec{\rho}$; then by Lemma 10.5 (the Δ -systems lemma) we can find a finite set r and an uncountable subset B of A such that r is the intersection of the supports of any two distinct elements of B . Fix $\delta < \gamma$ such that $r \subseteq \delta$; then if B were a down-antichain in P_γ it would easily follow that the restrictions of the elements of B to δ would be a down-antichain in P_δ . This contradicts the induction hypothesis that says P_δ is c.c.c. We conclude that P_γ is c.c.c. ■

Proposition 20.4. *Suppose $\alpha = \aleph_2^M$ and $M \models \text{“}P_\alpha \text{ is c.c.c.”}$ Let $A \in M$ satisfy $M \models \text{“}\text{card}(A) = \aleph_1^M\text{”}$, and let B be a subset of A in $M[G]$ for some generic filter G of P_α . Then in the notation of Proposition 20.2, B belongs to $M[G_\gamma]$ for some $\gamma < \alpha$.*

Proof. Let τ be a P_α -name for B . Working in $M[G]$, for each $x \in B$ find $\vec{\rho}_x \in G$ which forces $\check{x} \in \tau$. Since $\text{card}(B) \leq \aleph_1^M = \aleph_1^{M[G]}$ and the support of each $\vec{\rho}_x$ is finite, there exists $\gamma < \alpha = \aleph_2^M = \aleph_2^{M[G]}$ such that the support of each $\vec{\rho}_x$ is contained in γ . Then

$$B = \{x \in A : \text{some } \vec{\rho} \in G_\gamma \text{ extended by } \{1_\mu : \mu \geq \gamma\} \text{ forces } \check{x} \in \tau\}$$

belongs to $M[G_\gamma]$. ■

Exercises

(a) Prove Proposition 20.2.

(b) Let $\vec{p}, \vec{p}' \in P_\alpha$ and choose $\gamma < \alpha$ such that for every $\mu \geq \gamma$, either $\rho_\mu = 1'_\mu$ or $\rho'_\mu = 1'_\mu$. Prove that \vec{p} and \vec{p}' are incompatible in P_α if and only if $\vec{p}|_\gamma$ and $\vec{p}'|_\gamma$ are incompatible in P_γ .

21. Martin's axiom

Definition 21.1. Martin's axiom (MA) is the assertion that if P is a c.c.c. poset and $\{D_\alpha\}$ is a family of fewer than 2^{\aleph_0} dense subsets of P , then there is a filter G of P which intersects every D_α .

In general, one cannot hope to find a filter that intersects 2^{\aleph_0} dense sets. For example, let P be the infinite binary tree and for each branch ϕ let D_ϕ be the set of vertices not belonging to ϕ . Then each D_ϕ is dense but there is no filter of P that intersects every D_ϕ (exercise).

At the other extreme, it is a theorem of ZFC that for any countable family of dense subsets there is a filter that intersects them all (see Lemma 6.2). Thus, CH implies MA. But MA is also relatively consistent with \neg CH. Sometimes results proven using CH can actually be proven using only MA, and this implies that they are also relatively consistent with \neg CH. (For instance, this is true of the existence of pure states on $\mathcal{B}(H)$ which are not diagonalizable, proven in Theorem 13.5.)

In other cases, MA + \neg CH settles problems in a direction opposite to CH. So if some statement is proven relatively consistent using CH or \diamond , there is a reasonable chance that its negation can be proven relatively consistent using MA + \neg CH.

We will force the relative consistency of MA + $2^{\aleph_0} = \aleph_2$. The idea of the proof is to carry out an \aleph_2 -stage iterated forcing construction, where at each stage we add a filter which intersects \aleph_1 dense subsets of a given c.c.c. poset. It suffices to verify MA for posets of cardinality \aleph_1 , so we only consider posets whose underlying set is the ordinal \aleph_1 . Each stage of the construction adds at most \aleph_2 orderings of \aleph_1 , so a total of $\aleph_2^2 = \aleph_2$ orderings have to be handled and by arranging the construction carefully we are able to do this in \aleph_2 steps.

First we verify that it is enough to check MA for posets of cardinality \aleph_1 .

Lemma 21.2. *Assume $2^{\aleph_0} = \aleph_2$ and suppose that for any family of \aleph_1 dense subsets of any c.c.c. poset of cardinality \aleph_1 there is a filter that intersects them all. Then Martin's axiom is true.*

Proof. Let P be a c.c.c. poset and let $\{D_\alpha\}$ be a family of fewer than $2^{\aleph_0} = \aleph_2$ dense subsets of P . If this family is countable then there is a filter that intersects every D_α (this was noted above). Thus we may assume the family has cardinality \aleph_1 . If P is countable then let P' be the disjoint union of P and \aleph_1 , giving \aleph_1 its standard ordering and setting $p < q$ for every $p \in P$ and $q \in \aleph_1$. Then P' has cardinality \aleph_1 and the D_α are also dense subsets of P' , so by hypothesis there is a filter of P' that intersects every D_α , and its intersection with P is then a filter of P that intersects every D_α . So we may assume $\text{card}(P) \geq \aleph_1$.

We will find a sub-poset $Q \subseteq P$ whose cardinality is at most \aleph_1 , such that $D_\alpha \cap Q$ is dense in Q for all α and if $p, q \in Q$ have a common extension in P then they have a common extension in Q (this implies that Q is c.c.c.). Having done this, we can apply the above reduction to find a filter H of Q that intersects each D_α , and then $G = \{p \in P : p \geq q \text{ for some } q \in H\}$ is a filter of P that intersects each D_α , as desired.

We construct Q as follows. Let $f : P \times P \rightarrow P$ be any function such that $f(p, q)$ is a common extension of p and q provided such an extension exists (and otherwise $f(p, q)$ may be arbitrary). Also, for each α let $g_\alpha : P \rightarrow P$ be any function such that for all $p \in P$ we have $g_\alpha(p) \leq p$ and $g_\alpha(p) \in D_\alpha$. Now choose an element p_0 of P and let Q be the smallest subset of P which contains p_0 and is closed under f and the g_α . It is straightforward to check that $\text{card}(Q) \leq \aleph_1$ and that Q has the desired properties. ■

Theorem 21.3. *There is a notion of forcing P such that $M[G] \models \text{"MA} + 2^{\aleph_0} = \aleph_2\text{"}$, for any generic filter G of P .*

Proof. Working in M , let Q be the poset of bijections between subsets of $\mathcal{P}(\aleph_1)$ and \aleph_2 of size \aleph_1 . Then $M[H] \models "2^{\aleph_1} = \aleph_2"$, for any generic filter H of Q . The proof of this is analogous to the proof of Theorem 9.3. For the remainder of the proof we work in $M[H]$.

Let f be a function from \aleph_2 onto \aleph_2^2 such that the first coordinate of $f(\alpha)$ is at most α , for all $\alpha < \aleph_2$. We now define an \aleph_2 -stage finite support iterated forcing construction, such that each P_α is a c.c.c. preordering of \aleph_1 , as follows. Suppose P_β, π_β , etc., have been constructed for all $\beta < \alpha$. Then P_α is determined by the definition of finite support iterated forcing. We must now define π_α .

Recall the notion of nice names introduced in Lemma 17.6. For every subset of \aleph_1^2 introduced in a generic extension using the notion of forcing P_α , there is a nice P_α -name that evaluates to that subset. There are at most $(\aleph_1^{\aleph_0})^{\aleph_1} = \aleph_2$ nice names for subsets of \aleph_1^2 ; let $\{\sigma_\beta^\alpha : \beta < \aleph_2\}$ enumerate them. Say $f(\alpha) = \langle \gamma, \delta \rangle$. Since $\gamma \leq \alpha$, the P_γ -name σ_δ^γ has already been defined. This may not be a name for a c.c.c. preorder, but we can find a P_α -name π_α such that 1_α forces " π_α is a c.c.c. preordering of \aleph_1 with a greatest element, such that if σ_δ^γ is also a c.c.c. preordering of \aleph_1 with a greatest element then $\pi_\alpha = \sigma_\delta^\gamma$ ". This defines π_α , and the construction may proceed.

Let G be an $M[H]$ -generic filter of P_{\aleph_2} ; we claim that $M[H][G] \models "MA + 2^{\aleph_0} = \aleph_2"$. First, P_{\aleph_2} is c.c.c. by Proposition 20.2, so $\aleph_2^{M[H]} = \aleph_2^{M[H][G]}$. Work in $M[H][G]$. By a nice name argument similar to the one used above, there are at most \aleph_2 nice P_{\aleph_2} -names for subsets of \aleph_1 , so $2^{\aleph_0} \leq \aleph_2$. By Lemma 21.2, in order to verify MA it is enough to show that for any family of \aleph_1 dense subsets of any poset of cardinality \aleph_1 there is a filter that intersects them all. By Proposition 20.4, any ordering of \aleph_1 that appears in the final \aleph_2 -stage iterated forcing construction, and any family of \aleph_1 dense subsets of this ordering, already appear at some intermediate stage. Thus the ordering would have been represented by some nice P_γ -name σ_δ^γ , and there would have been some $\alpha < \aleph_2$ such that $\alpha \geq \gamma$ and $f(\alpha) = \langle \gamma, \delta \rangle$. We may also assume that the \aleph_1 dense subsets have appeared by stage α and the forcing at that stage would then have added a filter that intersects them all. Thus such a filter exists in the final generic extension.

We already showed that $2^{\aleph_0} \leq \aleph_2$. We cannot have $2^{\aleph_0} = \aleph_1$ because we have shown that for any \aleph_1 dense subsets of a c.c.c. poset there is a generic filter that intersects them all (see the comment immediately following Definition 21.1). So $2^{\aleph_0} = \aleph_2$.

Finally, let $P = Q * \pi$ where π is a Q -name for P_{\aleph_2} . ■

Exercises

- (a) In the example stated just after Definition 21.1, prove that no filter intersects every D_ϕ .
- (b) In the proof of Theorem 21.3, show that $M[H] \models "2^{\aleph_1} = \aleph_2"$.

22. Application: Suslin's problem, II

Recall (Definition 15.1) that a Suslin line is a totally ordered set that is dense, unbounded, complete, and c.c.c. but not order-isomorphic to \mathbb{R} , and Suslin's problem asks whether Suslin lines exist. Also recall that the existence of a Suslin line is equivalent to the existence of a Suslin tree (Definition 15.2), which is a tree of height \aleph_1 such that all branches and antichains are countable (Lemma 15.4).

We proved (Theorem 15.5) that the diamond principle implies the existence of Suslin lines. Now we show that Martin's axiom implies there are no Suslin lines.

Theorem 22.1. *Assume MA. Then Suslin lines do not exist.*

Proof. By Lemma 15.4, it will suffice to show that normal Suslin trees do not exist. Suppose to the contrary that T is a normal Suslin tree. Since all antichains are countable, T is c.c.c. For each $\alpha < \aleph_1$, let D_α be the set of all vertices of height at least α . Then each D_α is dense because every vertex has descendants at every level. By MA, there is a filter G of T that intersects each D_α . But then G is an uncountable branch, a contradiction. ■

We now give two other basic applications of Martin's axiom.

Theorem 22.2. *Assume MA. Then any subset of \mathbf{R} of cardinality $< 2^{\aleph_0}$ has measure zero.*

Proof. Let $S \subset \mathbf{R}$ and suppose $\text{card}(S) < 2^{\aleph_0}$. Fix $\epsilon > 0$ and let P be the set of all open subsets U of \mathbf{R} such that (1) U is a union of finitely many open intervals with rational endpoints and (2) the measure of U is less than ϵ . Order P by reverse inclusion; then P is c.c.c. because it is countable. Also, for each $r \in S$ let D_r be the set of $U \in P$ which contain r ; this is a dense subset of P since for any $U \in P$ an open interval I with rational endpoints that contains r can be found which is small enough that the measure of $U \cup I$ is still less than ϵ . By MA there is a generic filter of P which intersects every D_r . Its union is then an open set of measure at most ϵ which contains S . Since ϵ was arbitrary, we conclude that S has measure zero. ■

Recall that a subset of \mathbf{R} is meager if it is a countable union of nowhere-dense sets.

Theorem 22.3. *Assume MA and let $\{C_\alpha : \alpha < \kappa\}$ be a family of fewer than 2^{\aleph_0} meager subsets of \mathbf{R} . Then $\bigcup C_\alpha$ is meager.*

Proof. Let P be the set of all finite sequences of ordered pairs $\langle U, E \rangle$ such that (1) U is a finite union of open intervals with rational endpoints; (2) E is a finite subset of κ ; and (3) U is disjoint from $\bigcup_{\alpha \in E} C_\alpha$. Order P by setting $q \leq p$ if the sequence q is at least as long as the sequence p and for every $\langle U, E \rangle$ in p the corresponding $\langle U', E' \rangle$ in q satisfies $U \subseteq U'$ and $E \subseteq E'$.

P is c.c.c. because any uncountable subset of P contains an uncountable subset all of whose elements are sequences of the same length n , and this contains an uncountable subset all of whose elements involve the same sequence of U 's (since there are only countably sequences of U 's of length n); but any p and q of the same length with the same U 's are compatible.

For each $\alpha < \kappa$, the set D_α of $p \in P$ which contain an ordered pair $\langle U, E \rangle$ such that $\alpha \in E$ is dense. Also, for each open interval I with rational endpoints and each i the set $D_{I,i}$ of $p \in P$ whose i th ordered pair $\langle U, E \rangle$ satisfies $U \cap I \neq \emptyset$ is dense. There are κ dense sets of the first type and \aleph_0 dense sets of the second type, so by MA there is a filter G which intersects all of the above dense sets.

For each $i \in \mathbf{N}$ let V_i be the union of all U such that $\langle U, E \rangle$ is the i th ordered pair in some $p \in G$ (for some E). Since $G \cap D_{I,i} \neq \emptyset$ for all I , it follows that V_i is a dense open subset of \mathbf{R} . Also, since $G \cap D_\alpha \neq \emptyset$, for each $\alpha < \kappa$ we can find $\langle U, E \rangle \in p \in G$ such that $\alpha \in E$; if this is the i th pair in G then $V_i \cap A_\alpha = \emptyset$. So $\bigcap V_i$ is a countable intersection of dense open sets whose complement contains every C_α . Thus $\bigcup C_\alpha$ is meager. ■

Exercises

(a) Assuming $\text{MA} + \neg \text{CH}$, prove that any union of \aleph_1 measure zero subsets of \mathbf{R} is a measure zero subset of \mathbf{R} .

23. Application: The Whitehead problem, II

Now we show that $\text{MA} + \neg \text{CH}$ implies that there is a counterexample to Whitehead's problem, i.e., an abelian group with only trivial extensions by \mathbf{Z} that is not free. This proof is also due to Shelah.

Shelah's example is defined as follows. It has generators a_α for all $\alpha < \aleph_1$ and b_α^n for every limit ordinal $\alpha < \aleph_1$ and every $n \in \mathbf{N}$. For each limit ordinal α choose a sequence (α_n) of ordinals which increase to α and include the relation

$$b_\alpha^n = a_{\alpha_n} + 2b_{\alpha_n}^{n+1}$$

for all $n \in \mathbf{N}$. Let G be the abelian group defined by these generators and relations.

For each $\alpha < \aleph_1$ let G_α be the subgroup of G generated by all a_β with $\beta < \alpha$ and all b_β^n with $\beta \leq \alpha$ a limit ordinal and $n \in \mathbf{N}$. Note that $\bigcup_{\beta < \alpha} G_\beta$ does not equal G_α when α is a limit ordinal.

Lemma 23.1. *G is not free.*

Proof. Suppose G is free and let \mathcal{B} be a basis. Inductively define a sequence of ordinals α_n as follows. Let $\alpha_0 = 0$. Given α_n , G_{α_n} is a countable subgroup of G so it is contained in the span of a countable subset B' of \mathcal{B} . Find α_{n+1} such that $G_{\alpha_{n+1}}$ contains B' .

Let $\alpha = \sup \alpha_n$ and let $G'_\alpha = \bigcup_n G_{\alpha_n}$. Then G'_α is spanned by $G'_\alpha \cap \mathcal{B}$, so $G = G'_\alpha \oplus G/G'_\alpha$ and both summands are free. However, this is impossible because the nonzero element $b_\alpha^0 + G'_\alpha$ in G/G'_α satisfies

$$b_\alpha^0 + G'_\alpha = 2b_\alpha^1 + G'_\alpha = 4b_\alpha^2 + G'_\alpha = \cdots,$$

contradicting the fact that G/G'_α is free. We conclude that G is not free. ■

Lemma 23.2. G_α/G_β is free for all $\beta < \alpha < \aleph_1$.

Proof. Exercise. ■

A subgroup H of a torsion-free group is pure if $na \in H$ implies $a \in H$, for any $n \in \mathbb{N}$. We need the following fact from group theory: any subgroup of a free abelian group is free.

Let A be an abelian group and let $\pi : A \rightarrow G$ be a surjective homomorphism with kernel \mathbf{Z} . Define P to be the poset of all homomorphisms $\rho : H \rightarrow A$ such that H is a finitely generated pure subgroup of G and $\pi \circ \rho = \text{id}_H$, with $\rho \leq \rho'$ if ρ is an extension of ρ' .

Lemma 23.3. P is c.c.c.

Proof. Let $S \subset P$ be uncountable. We claim that there is a pure free subgroup H of G which contains the domains of uncountably many $\rho \in S$. This is enough because if \mathcal{B} is a basis for H , then the domain of any ρ contained in H is contained in the span of a finite subset \mathcal{B}_ρ of \mathcal{B} . By the Δ -systems lemma (Lemma 10.5) there is an uncountable family S' of $\rho \in S$ and a finite set $\mathcal{B}_0 \subseteq \mathcal{B}$ such that $\mathcal{B}_\rho \cap \mathcal{B}_{\rho'} = \mathcal{B}_0$ for all distinct $\rho, \rho' \in S'$. For each $\rho \in S'$ let $\bar{\rho}$ be an extension of ρ to the span of \mathcal{B}_ρ (this is possible because the domain of ρ is pure, so that $\text{span}(\mathcal{B}_\rho)/\text{dom}(\rho)$ is free by the structure theorem for finitely generated abelian groups). Then there are only countably many possible functions $\bar{\rho}|_{\mathcal{B}_0}$, so some distinct $\bar{\rho}, \bar{\rho}' \in S'$ have the same restriction to \mathcal{B}_0 and hence they have a common extension to the span of $\mathcal{B}_\rho \cup \mathcal{B}_{\rho'}$. Thus ρ and ρ' are compatible, and we conclude that P is c.c.c.

To prove the claim, find $n \in \mathbb{N}$ such that the domains of uncountably many $\rho \in S$ have bases of size n . Let S' be the set of all such ρ . Then let H_0 be a maximal pure subgroup of G with the property that $H_0 \subseteq \text{dom}(\rho)$ for uncountably many $\rho \in S'$. Let T be the set of all $\rho \in S'$ whose domain contains H_0 .

Recursively define homomorphisms $\rho_\alpha \in T$ and countable free pure subgroups H_α of G for $\alpha < \aleph_1$ as follows. At successor stages, suppose H_α has been defined and let $H_\alpha \subseteq G_{\alpha'}$. For all but countably many $\rho \in T$ we have $\text{dom}(\rho) \cap G_{\alpha'} = H_0$. (Otherwise, since $G_{\alpha'}$ is countable, some $a \in G_{\alpha'} - H_0$ would be contained in the domains of uncountably many $\rho \in T$ and this would contradict maximality of H_0 .) So choose $\rho_\alpha \in T$ distinct from all ρ_β for $\beta < \alpha$ such that $\text{dom}(\rho_\alpha) \cap G_{\alpha'} = H_0$. Also let $H_{\alpha+1}$ be the smallest pure subgroup of G that contains H_α and $\text{dom}(\rho_\alpha)$. If $H_{\alpha+1} \subseteq G_{\alpha''}$ then $H_{\alpha+1}/H_\alpha$ is isomorphic to a subgroup of $G_{\alpha''}/G_{\alpha'}$, which is free, so that freeness of H_α implies that $H_{\alpha+1}$ is also free.

At limit stages let $H_\alpha = \bigcup_{\beta < \alpha} H_\beta$. Since $H_{\beta+1}/H_\beta$ is free for all $\beta < \alpha$, it follows that H_α is free just as in the proof of Theorem 18.8. Also, purity of H_α trivially follows from purity of all H_β for $\beta < \alpha$.

Finally, let $H = \bigcup_{\alpha < \aleph_1} H_\alpha$. Then H is free and pure just as above. It contains $\text{dom}(\rho_\alpha)$ for all $\alpha < \aleph_1$, so the claim is proven. ■

Theorem 23.4. Assume $MA + \neg CH$. Then G is a group of cardinality \aleph_1 that is not free and all of whose extensions by \mathbf{Z} are trivial.

Proof. G is not free by Lemma 23.1. Let $\pi : A \rightarrow G$ be an extension by \mathbf{Z} and let P be the poset defined just before Lemma 23.3. We must show that π splits.

For each $a \in G$ let D_a be the set of all $\rho \in P$ whose domain contains a . We claim that D_a is dense. To see this, let $\rho \in P$ and let G_α contain a and the domain of ρ . Let \mathcal{B} be a basis for G_α and let \mathcal{B}_0 be a finite subset of \mathcal{B} whose span contains a and the domain of ρ . Let H be the span of \mathcal{B}_0 ; then $H/\text{dom}(\rho)$ is free, so we can choose a basis and extend ρ to a splitting of $\pi|_{\pi^{-1}(H)} : \pi^{-1}(H) \rightarrow H$ one basis element at a time (cf. the proof of Proposition 18.2). This shows that ρ has an extension in D_a .

Since $\text{card}(G) = \aleph_1 < 2^{\aleph_0}$, Martin's axiom implies that there is a filter F of P that intersects every D_a . The union of F is then a homomorphism from G to A that splits π , as desired. ■

Exercises

(a) Prove Lemma 23.2.

References

D. H. Fremlin, *Consequences of Martin's Axiom*, 1984.