

Ultrafilters and Partial Products of Infinite Cyclic Groups

Andreas Blass*

Mathematics Department
University of Michigan
Ann Arbor, MI 48109–1109, U.S.A.
ablass@umich.edu

Saharon Shelah†

Mathematics Department
Hebrew University
Jerusalem 91904, Israel
and
Mathematics Department
Rutgers University
New Brunswick, NJ 08903, U.S.A.

Abstract

We consider, for infinite cardinals κ and $\alpha \leq \kappa^+$, the group $\Pi(\kappa, <\alpha)$ of sequences of integers, of length κ , with non-zero entries in fewer than α positions. Our main result tells when $\Pi(\kappa, <\alpha)$ can be embedded in $\Pi(\lambda, <\beta)$. The proof involves some set-theoretic results, one about families of finite sets and one about families of ultrafilters.

*Partially supported by NSF grant DMS-0070723. Part of this paper was written during a visit of the first author to the Centre de Recerca Matemàtica in Barcelona.

†Partially supported by US-Israel Binational Science Foundation grant 2002323. Publication number 854 of the second author.

1 Introduction

For an infinite cardinal κ , let \mathbb{Z}^κ be the direct product of κ copies of the additive group \mathbb{Z} of integers. An element of \mathbb{Z}^κ is thus a function¹ $x : \kappa \rightarrow \mathbb{Z}$, and we define its *support* to be the set

$$\text{supp}(x) = \{\xi \in \kappa : x(\xi) \neq 0\}.$$

The partial products mentioned in the title of this paper are the subgroups of \mathbb{Z}^κ of the form

$$\Pi(\kappa, <\alpha) = \{x \in \mathbb{Z}^\kappa : |\text{supp}(x)| < \alpha\}$$

where α is an infinite cardinal no larger than the successor cardinal κ^+ of κ . Notice that $\Pi(\kappa, <\kappa^+)$ is the full product \mathbb{Z}^κ . At the other extreme, $\Pi(\kappa, <\omega)$ is the direct sum of κ copies of \mathbb{Z} , i.e., the free abelian group generated by the κ standard unit vectors e_ξ defined by $e_\xi(\xi) = 1$ and $e_\xi(\eta) = 0$ for $\xi \neq \eta$.

The main result in this paper gives necessary and sufficient conditions for one partial product of \mathbb{Z} 's to be isomorphically embeddable in another.

Theorem 1 *$\Pi(\kappa, <\alpha)$ is isomorphic to a subgroup of $\Pi(\lambda, <\beta)$ if and only if either*

1. $\kappa \leq \lambda$ and $\alpha \leq \beta$ or
2. $\kappa \leq \lambda^{<\beta}$ and $\alpha = \omega$.

Part of this was proved in [2, Theorem 23 and Remark 28], using a well-known result from set theory, the Δ -system lemma. Specifically, the results in [2] establish Theorem 1 when α and β are regular, uncountable cardinals smaller than all measurable cardinals. In the present paper, we complete the proof by handling the cases of singular cardinals and cardinals above a measurable one. In contrast to the situation in [2], this will involve developing some new results in set theory, rather than only invoking classical facts.

The set theoretic facts we need are the following two.

Theorem 2 *Let κ be an infinite cardinal and let $(F_\xi)_{\xi \in \kappa}$ be a κ -indexed family of nonempty, finite sets.*

¹We use the standard notational conventions of set theory, whereby a cardinal number is an initial ordinal number and is identified with the set of all smaller ordinals. In particular, the cardinal of countable infinity is identified with the set ω of natural numbers.

1. *There exists a set X such that*

$$|\{\xi \in \kappa : |F_\xi \cap X| = 1\}| = \kappa. \quad (1)$$

2. *The set X in (1) can be chosen so that*

- *$|X|$ has cardinality 1 or $\text{cf}(\kappa)$ or κ ,*
- *every subset of X with the same cardinality as X has the property in (1), and*
- *each element of X is the unique element of $F_\xi \cap X$ for at least one ξ .*

Theorem 3 *Let κ be an infinite cardinal and let $(\mathcal{U}_\xi)_{\xi \in \kappa}$ be a κ -indexed family of non-principal ultrafilters on κ . Then there exists $X \subseteq \kappa$ such that $|X| = \kappa$ and, for each $\xi \in X$, $X \notin \mathcal{U}_\xi$.*

We prove these two set-theoretic theorems in Section 2. Then, in Section 3, we apply them to prove Theorem 1.

2 Set Theory

We thank Stevo Todorčević for suggesting a simplification, using the Δ -system lemma, of our original proof of Theorem 2. That suggestion led us, by further simplification, to the following proof, which doesn't need the Δ -system lemma.

Proof of Theorem 2 Let κ and $(F_\xi)_{\xi \in \kappa}$ be given, as in the hypothesis of the theorem. If there exists some x that lies in F_ξ for κ values of ξ , then $X = \{x\}$ obviously satisfies the conclusion of the theorem. So we assume from now on that each x lies in F_ξ for fewer than κ values of ξ .

Lemma 4 *There exists a set $D \subseteq \kappa$ with $|D| = \kappa$ and there is a function assigning, to each $\xi \in D$, some $x_\xi \in F_\xi$ with the following property. Whenever $x_\alpha \in F_\beta$ with $\alpha, \beta \in D$, then $x_\alpha = x_\beta$.*

Proof We begin by simplifying a special case that would otherwise interfere with the main argument. The special case is that κ is singular, say with cofinality μ , and that there are, for arbitrarily large $\lambda < \kappa$, finite sets G such

that $|\{\xi < \kappa : F_\xi = G\}| = \lambda$. That is, there are finite sets G that are repeated nearly κ times in the family $(F_\xi)_{\xi \in \kappa}$. (Note that no set can be repeated κ times, thanks to our standing assumption that no x occurs in F_ξ for κ values of ξ .) In this case, we can fix an increasing μ -sequence of cardinals λ_i with supremum κ , and we can fix finite sets G_i such that each G_i is equal to F_ξ for λ_i values of ξ . Then we apply the argument given below to the family $(G_i)_{i \in \mu}$ instead of the original family $(F_\xi)_{\xi \in \kappa}$. The result will be a set $D' \subseteq \mu$ of cardinality μ and a function assigning to each $i \in D'$ some $x'_i \in G_i$ such that, whenever $x'_i \in G_j$ with $i, j \in D'$ then $x'_i = x'_j$. Then we define D to be the set of all those ξ such that $F_\xi = G_i$ for some $i \in D'$, and we define, for each such ξ , x_ξ to be x'_i , where $i \in D'$ with $x_\xi = x'_i$. (The defining property of D' and the x'_i 's ensures that this x'_i is uniquely determined for each ξ .) It is easy to verify that D and the x_ξ 's are as required by the lemma.

This completes the proof in the exceptional case, so we assume from now on that its case hypothesis does not hold. This implies that, for any set A of cardinality $< \kappa$, the number of ξ for which $F_\xi \subseteq A$ is also $< \kappa$. Indeed, since A has fewer than κ distinct finite subsets G , the number of ξ such that $F_\xi \subseteq A$ is the sum, over these G , of their multiplicities in the sequence $(F_\xi)_{\xi \in \kappa}$. These multiplicities are all $< \kappa$, so the only way their sum, over the fewer than κ G 's, can be κ is for the hypothesis of the exceptional case to hold.

We are now ready to start building the required D and the required function $\xi \mapsto x_\xi$ inductively. We begin with D empty, and we enlarge it step by step, stopping when its cardinality reaches κ . At each step, we shall choose a suitable x and add to D all those ξ such that $x \in F_\xi$; for each of these ξ , we shall set $x_\xi = x$. In order for this definition to be consistent and to satisfy the requirements of the lemma, our choice of x is subject to several constraints:

- x is not in F_η for any η previously put into D .
- No F_ξ contains both x and any x_η for η previously put into D .
- x is in F_ξ for some ξ .

The first of these constraints ensures that the requirement in the lemma is satisfied when α is one of the ξ 's being added at the current step and β was put into D earlier. The second ensures the requirement of the lemma when β is one of the ξ 's being added at the current step and α was put into D earlier.

(In both cases, we ensure that $x_\alpha \notin F_\beta$.) The requirement of the lemma will also hold when both α and β are among the currently added ξ 's, because then $x_\alpha = x_\beta = x$. The third constraint merely ensures that D acquires at least one new element per step; any ξ as in the third constraint is put into D , and it wasn't previously in D because of the first constraint.

To complete the proof of the lemma, we must show that, as long as $|D| < \kappa$, we can find an x satisfying all the constraints.

In fact, the second constraint is redundant. If F_ξ and η violated it, then ξ would have been put into D already at the same step where η was added, because we always add all F 's that contain the currently chosen x . Thus, the first constraint would be violated with ξ in the role of η . So we need only show that, when $|D| < \kappa$, we can choose x so as to satisfy the first and third constraints. The union of the F_η 's for η previously put into D is a set A of cardinality $< \kappa$, because $|D| < \kappa$ and the F_η 's are finite. We saw above that such an A cannot include F_ξ for κ values of ξ . So we can choose a $\xi < \kappa$ with $F_\xi \not\subseteq A$ and we can choose $x \in F_\xi - A$. This x clearly satisfies the first and third constraints, so the proof of the lemma is complete. \square

Fix D and $\xi \mapsto x_\xi$ as in the lemma. We next normalize the D a bit as follows. Let \sim be the equivalence relation on D defined by

$$\xi \sim \eta \iff x_\xi = x_\eta.$$

We shall arrange that one of the following three alternatives holds.

1. D is a single equivalence class, i.e., all the x_ξ are equal.
2. Each equivalence class is a singleton, i.e., all the x_ξ are distinct.
3. κ is singular, the number of equivalence classes is $\mu = \text{cf}(\kappa)$, and their sizes form a cofinal subset of κ of order-type μ .

We can arrange this simply by shrinking D (while keeping its cardinality equal to κ of course). If there is an equivalence class of size κ , then replacing D by that equivalence class attains alternative (1). If there are κ equivalence classes, then replacing D by a selector attains alternative (2). So we may assume that there are $< \kappa$ equivalence classes, each of size $< \kappa$. Thus, κ is singular; let μ be its cofinality. The sizes of the equivalence classes must be unbounded below κ , for otherwise their union would be smaller than κ (being at most the bound times μ). So we can choose a μ -sequence of equivalence

classes of increasing cardinalities approaching κ . Replacing D by the union of these classes attains alternative (3).

Finally, we let $X = \{x_\xi : \xi \in D\}$ and we check that it has the properties required in the theorem. If $\beta \in D$, then $x_\beta \in F_\beta$ and, by the requirement in the lemma, no $x_\alpha \neq x_\beta$ can be in F_β . So $|F_\beta \cap X| = 1$ for all $\beta \in D$. Since $|D| = \kappa$, part 1 of the theorem is satisfied. The cardinality of X is the number of equivalence classes with respect to \sim in D , and our normalization of D ensures that this is 1 or κ or $\text{cf}(\kappa)$. The normalization also ensures that any subset of X of the same cardinality as X arises from a subset of D that shares the properties we obtained for D . So any such subset also works in part 1 of the theorem. Finally, each element $x \in X$ is of the form x_ξ for some $\xi \in D$ and therefore is, thanks to the requirement on D in the lemma, the unique element of $F_\xi \cap X$. \square

Proof of Theorem 3 Let κ and $(\mathcal{U}_\xi)_{\xi \in \kappa}$ be as in the hypothesis of the theorem. Partition κ into κ sets A_μ (with $\mu \in \kappa$), each of cardinality κ . If one of these A_μ can serve as X in the conclusion of the theorem, then nothing more needs to be done. So assume that this is not the case, i.e., assume that, for each μ , there is some $\xi(\mu) \in A_\mu$ such that $A_\mu \in \mathcal{U}_{\xi(\mu)}$. Being non-principal, $\mathcal{U}_{\xi(\mu)}$ also contains $A_\mu - \{\xi(\mu)\}$.

Let $X = \{\xi(\mu) : \mu \in \kappa\}$. For each element of X , say $\xi(\mu)$, we have seen that $\mathcal{U}_{\xi(\mu)}$ contains a set disjoint from X , namely $A_\mu - \{\xi(\mu)\}$. Therefore $X \notin \mathcal{U}_{\xi(\mu)}$, and the proof is complete. \square

3 Proof of Theorem 1

We begin by showing that, if one of the cardinality conditions 1 and 2 in Theorem 1 is satisfied, then we can embed $\Pi(\kappa, <\alpha)$ in $\Pi(\lambda, <\beta)$.

If $\kappa \leq \lambda$, then we can embed \mathbb{Z}^κ into \mathbb{Z}^λ by extending any κ -sequence $x \in \mathbb{Z}^\kappa$ by zeros to have length λ . This does not alter the support, so it embeds $\Pi(\kappa, <\alpha)$ into $\Pi(\lambda, <\beta)$ (as a pure subgroup) for any $\beta \geq \alpha$.

This completes the proof if condition 1 in the theorem is satisfied. If condition 2 is satisfied, then, since $\alpha = \omega$, the group $\Pi(\kappa, <\alpha)$ is a free abelian group of rank $\kappa \leq \lambda^{<\beta}$. Since $\Pi(\lambda, <\beta)$ has cardinality $\lambda^{<\beta}$, its rank is also $\lambda^{<\beta}$. (The only way for a torsion-free abelian group to have rank different from its cardinality is to have finite rank, which is clearly not the case for $\Pi(\lambda, <\beta)$.) So it has a free subgroup of rank $\lambda^{<\beta}$, and we have the required embedding.

Remark 5 Nöbeling proved in [6] that the subgroup of \mathbb{Z}^λ consisting of the bounded functions is a free abelian group. Intersecting it with $\Pi(\lambda, <\beta)$, we get a pure free subgroup of $\Pi(\lambda, <\beta)$ of rank $\lambda^{<\beta}$. Thus, under condition 2 of the theorem, we get an embedding of $\Pi(\kappa, <\alpha)$ into $\Pi(\lambda, <\beta)$ as a pure subgroup. Therefore, Theorem 1 would remain correct if we replaced “subgroup” with “pure subgroup.”

We now turn to the more difficult half of Theorem 1, assuming the existence of the embedding of groups and deducing one of the cardinality conditions. Since $\Pi(\lambda, <\beta)$ has cardinality $\lambda^{<\beta}$ and $\Pi(\kappa, <\alpha)$ has cardinality at least κ , the existence of an embedding of the latter into the former obviously implies that $\kappa \leq \lambda^{<\beta}$. So if $\alpha = \omega$ then we have condition 2 of the theorem. Therefore, we assume from now on that α is uncountable; our goal is to deduce condition 1.

For this purpose, we need to assemble some information about the given embedding $j : \Pi(\kappa, <\alpha) \rightarrow \Pi(\lambda, <\beta)$. The embedding is, of course, determined by its λ components, i.e., its compositions with the λ projection functions $p_\nu : \Pi(\lambda, <\beta) \rightarrow \mathbb{Z}$. (Here and in all that follows, the variable ν is used for elements of λ .) We write j_ν for $p_\nu \circ j : \Pi(\kappa, <\alpha) \rightarrow \mathbb{Z}$. Thus, for any $x \in \Pi(\kappa, <\alpha)$, $j_\nu(x)$ is the ν^{th} component of the λ -sequence $j(x)$.

The structure of homomorphisms, like j_ν , from $\Pi(\kappa, <\alpha)$ to \mathbb{Z} can be determined, thanks to the following theorem of Balcerzyk [1]. (This theorem extends earlier results of Specker [7] for $\kappa = \omega$ and Łoś (see [4, Theorem 94.4]) for κ smaller than all measurable cardinals; it was in turn extended by Eda [3] to allow arbitrary slender groups in place of \mathbb{Z} .) To state it, we need one piece of notation.

If \mathcal{U} is a countably complete ultrafilter on a set A and if x is any function from A to a countable set (such as \mathbb{Z}), then x is constant on some set in \mathcal{U} , and we denote that constant value by $\mathcal{U}\text{-lim } x$.

Theorem 6 (Balcerzyk) *Let A be any set and let $h : \mathbb{Z}^A \rightarrow \mathbb{Z}$ be a homomorphism. Then there exist finitely many countably complete ultrafilters \mathcal{U}_i on A and there exist integers c_i (indexed by the same finitely many i 's) such that, for all $x \in \mathbb{Z}^A$,*

$$h(x) = \sum_i c_i \cdot \mathcal{U}_i\text{-lim } x.$$

We shall refer to the sum in this theorem as the *Balcerzyk formula* for h . Whenever it is convenient, we shall assume that, in a Balcerzyk formula, all

the \mathcal{U}_i are distinct and all the c_i are non-zero. This can be arranged simply by combining any terms that involve the same ultrafilter and omitting any terms with zero coefficients.

The theorem easily implies that the group of homomorphisms from \mathbb{Z}^A to \mathbb{Z} is freely generated by the homomorphisms \mathcal{U} -lim for countably complete ultrafilters \mathcal{U} on A .

Notice that among the countably complete ultrafilters are the principal ultrafilters, and that the homomorphism \mathcal{U} -lim associated to the principal ultrafilter \mathcal{U} at some $a \in A$ is simply the projection $p_a : \mathbb{Z}^A \rightarrow \mathbb{Z} : x \rightarrow x(a)$. If $|A|$ is smaller than all measurable cardinals, then the principal ultrafilters are the only countably complete ultrafilters on A , so homomorphisms from \mathbb{Z}^A to \mathbb{Z} are simply finite linear combinations of projections.

Corollary 7 *If $h : \mathbb{Z}^A \rightarrow \mathbb{Z}$ is a homomorphism, then there are only finitely many $a \in A$ such that the standard unit vector e_a is mapped to a non-zero value by h .*

Proof For $h(e_a)$ to be non-zero, one of the \mathcal{U}_i in the theorem must be the principal ultrafilter at a . \square

We wish to apply this information to the homomorphisms j_ν , whose domain is only $\Pi(\kappa, <\alpha)$, not all of \mathbb{Z}^κ . Fortunately, the preceding corollary carries over to the desired context, thanks to our assumption above that α is uncountable.

Corollary 8 *For each $\nu \in \lambda$, there are only finitely many $\xi \in \kappa$ such that $j_\nu(e_\xi) \neq 0$.*

Proof Suppose not. Then there is a countably infinite set $A \subseteq \kappa$ such that, for each $\xi \in A$, $j_\nu(e_\xi) \neq 0$. View \mathbb{Z}^A as a subgroup of \mathbb{Z}^κ , simply by extending functions by 0 on $\kappa - A$. Since α is uncountable, we have made \mathbb{Z}^A a subgroup of $\Pi(\kappa, <\alpha)$, the domain of j_ν . So we can apply Corollary 7 to (the restriction to \mathbb{Z}^A of) j_ν and conclude that $j_\nu(e_\xi) \neq 0$ for only finitely many $\xi \in A$. This contradicts our choice of A . \square

For each $\nu \in \lambda$, let

$$F_\nu = \{\xi \in \kappa : j_\nu(e_\xi) \neq 0\}.$$

So each F_ν is finite. On the other hand, since j is an embedding, we have, for each $\xi \in \kappa$, that $j(e_\xi) \neq 0$ and therefore $\xi \in F_\nu$ for at least one $\nu \in \lambda$.

Thus, κ is the union of the λ finite sets F_ν , which implies that $\kappa \leq \lambda$. This proves the first part of condition 1 of the theorem.

Before turning to the second part, we note, since we shall need it later, that the preceding argument shows not only that $\kappa \leq \lambda$ but that

$$\kappa \leq |\{\nu \in \lambda : F_\nu \neq \emptyset\}|.$$

To complete the proof of condition 1 of the theorem, it remains to show that $\alpha \leq \beta$. Suppose, toward a contradiction, that $\beta < \alpha$. So $\beta^+ \leq \alpha \leq \kappa^+$ and therefore $\beta \leq \kappa$. Therefore (by the first part of this proof), $\mathbb{Z}^\beta = \Pi(\beta, <\beta^+)$ embeds in $\Pi(\kappa, <\alpha)$, which in turn embeds in $\Pi(\lambda, <\beta)$. So instead of dealing with an embedding $\Pi(\kappa, <\alpha) \rightarrow \Pi(\lambda, <\beta)$, we can deal with an embedding $j : \mathbb{Z}^\beta \rightarrow \Pi(\lambda, <\beta)$. In other words, we can assume, without loss of generality, that $\kappa = \beta$ and $\alpha = \beta^+$.

We record for future reference that we have already reached a contradiction if $\beta = \omega$, for then $\Pi(\lambda, <\beta)$ is the free abelian group on λ generators while, by a theorem of Specker [7], \mathbb{Z}^β is not free. So the latter cannot be embedded into the former. Thus, we may assume, for the rest of this proof, that β is uncountable.

As before, we write j_ν for the homomorphism $\mathbb{Z}^\beta \rightarrow \mathbb{Z}$ given by the ν^{th} component of j , for each $\nu \in \lambda$. Also as before, we write F_ν for the set of $\xi \in \beta$ such that $j_\nu(e_\xi) \neq 0$. It will be useful to write the Balcerzyk formula for j_ν with the principal and non-principal ultrafilters separated. Note that the principal ultrafilters that occur here are concentrated at the points of F_ν . Thus, we have

$$j_\nu(x) = \sum_{\xi \in F_\nu} a_\xi^\nu \cdot x(\xi) + \sum_{\mathcal{U} \in \mathbb{U}_\nu} b_\mathcal{U}^\nu \cdot \mathcal{U}\text{-lim } x \quad (2)$$

where \mathbb{U}_ν is a finite set of non-principal, countably complete ultrafilters on β . As before, we assume, without loss of generality, that all the a and b coefficients are non-zero.

We recall that we showed, in the proof of $\kappa \leq \lambda$, that $F_\nu \neq \emptyset$ for at least β values of ν (since the κ of that proof is now equal to β). So we can apply Theorem 2 to find an $X \subseteq \beta$ with the following properties.

1. There are β values of ν , which we call the *special* values, such that $X \cap F_\nu$ is a singleton.
2. $|X|$ is one of 1, $\text{cf}(\beta)$, and β .

3. Every subset of X of the same cardinality as X shares with X the property in item 1 above.
4. Each $\xi \in X$ is, for at least one ν , the unique element of $X \cap F_\nu$.

It will be useful to select, for each $\xi \in X$, one ν as in item 4 and to call it $\nu(\xi)$. Notice that $\nu(\xi)$ is always special (as defined in item 1).

In the course of the proof, we will occasionally replace X by a subset of the same cardinality, relying on property 3 of X to ensure that all the properties listed for X remain correct for the new X . To avoid an excess of subscripts, we will not give these X 's different names. Rather, at each stage of the proof, X will refer to the current set, which may be a proper subset of the original X introduced above.

The basic idea of the proof is quite simple, so we present it first and afterward indicate how to handle all the issues that arise in its application.

Consider any $x \in \mathbb{Z}^\beta$ whose support is exactly X . Then for each special ν the first sum in (2) reduces to a single term, because exactly one $\xi \in F_\nu$ has $x(\xi) \neq 0$. So this formula reads

$$j_\nu(x) = a_\xi^\nu \cdot x(\xi) + \sum_{\mathcal{U} \in \mathbb{U}_\nu} b_\mathcal{U}^\nu \cdot \mathcal{U}\text{-lim } x \quad (3)$$

where ξ is the unique element of $X \cap F_\nu$. If we knew that none of the ultrafilters $\mathcal{U} \in \mathbb{U}_\nu$ contain X , then all the corresponding limits $\mathcal{U}\text{-lim } x$ would vanish, since \mathcal{U} contains a set (namely the complement of X) on which x is identically 0. In this case, we would have

$$j_\nu(x) = a_\xi^\nu \cdot x(\xi) \neq 0.$$

If this happened for β distinct values of ν , then all these values would be in the support of $j(x)$, contradicting the fact that $j(x) \in \Pi(\lambda, <\beta)$.

This is the basic idea; the rest of the proof is concerned with the obvious difficulty that we do not immediately have β values of ν for which the ultrafilters $\mathcal{U} \in \mathbb{U}_\nu$ do not contain X .

Of course, this difficulty cannot arise if $|X| = 1$, as the ultrafilters in question are non-principal. So the proof is complete if there is some ξ that lies in β of the sets F_ν , for then $\{\xi\}$ could serve as X . From now on, we assume that there is no such ξ .

More generally, the difficulty cannot arise, and so the proof is complete, if $|X|$ is smaller than all measurable cardinals, because then there are no non-principal, countably complete ultrafilters to contribute to the second sum in (2). So we may assume that there is at least one measurable cardinal $\leq |X|$.

There remain the cases that $|X| = \beta$ and that $|X| = \text{cf}(\beta) < \beta$. It turns out to be necessary to subdivide the former case according to whether $\text{cf}(\beta) = \omega$ or not. We handle the three resulting cases in turn.

Case 1: $|X| = \beta$ and $\text{cf}(\beta) > \omega$.

Recall that we chose, for each $\xi \in X$, some $\nu(\xi)$ such that $X \cap F_{\nu(\xi)} = \{\xi\}$. Thus, equation (3) holds when we put $\nu(\xi)$ in place of ν .

There are only countably many possible values for $|\mathbb{U}_{\nu(\xi)}|$ because these cardinals are finite. Since $|X|$ has, by the case hypothesis, uncountable cofinality, X must have a subset, of the same cardinality β , such that $|\mathbb{U}_{\nu(\xi)}|$ has the same value, say l , for all ξ in this subset. Replace X with this subset; as remarked above, we do not, with this replacement, lose any of the properties of X listed above. Now we can, for each ξ in (the new) X , enumerate $\mathbb{U}_{\nu(\xi)}$ as $\{\mathcal{U}_k(\xi) : k < l\}$.

Next, apply Theorem 3 l times in succession, starting with the current X . At step k (where $0 \leq k < l$), replace the then current X with a subset, still of cardinality β , such that, for each ξ in (the new) X , $\mathcal{U}_k(\xi)$ does not contain X . Thus, for the final X , after these l shrinkings, we have that, for all $\xi \in X$, and all $\mathcal{U} \in \mathbb{U}_{\nu(\xi)}$, $X \notin \mathcal{U}$. This is exactly what we need in order to apply the basic idea, explained above, to all the ν 's of the form $\nu(\xi)$ for $\xi \in X$. Since the function $\xi \mapsto \nu(\xi)$ is obviously one-to-one, there are β of these ν 's, and so we have the required contradiction.

Notice that the case hypothesis that β has uncountable cofinality was used in order to get a single cardinal l for $|\mathbb{U}_{\nu(\xi)}|$, independent of ξ , which was used in turn to fix the number of subsequent shrinkings of X . Without a fixed l , there would be no guarantee of a final X to which the basic idea can be applied. This is why the following case must be treated separately. It is the only case where the actual values of x , not just its support, will matter.

Case 2: $|X| = \beta$ and $\text{cf}(\beta) = \omega$.

Recall that we have already obtained a contradiction when $\beta = \omega$, so in the present case β is a singular cardinal. Fix an increasing ω -sequence $(\beta_n)_{n \in \omega}$ of uncountable regular cardinals with supremum β . Partition X into countably many sets X_n with $|X_n| = \beta_n$. As in the proof of Case 1, we can shrink each X_n , without decreasing its cardinality, so that:

- The cardinality of $\mathbb{U}_{\nu(\xi)}$ depends only on n , not on the choice of $\xi \in X_n$; call this cardinality $l(n)$.
- For all $\xi \in X_n$, no ultrafilter in $\mathbb{U}_{\nu(\xi)}$ contains X_n .

Here and below, when we shrink the X_n 's, it is to be understood that X is also shrunk, to the union of the new X_n 's. As long as the cardinality of each X_n remains β_n , the cardinality of X remains β .

As before, we use the notation $\{\mathcal{U}_k(\xi) : k < l(n)\}$ for an enumeration of $\mathbb{U}_{\nu(\xi)}$ when $\xi \in X_n$.

Notice that each $\mathcal{U}_k(\xi)$, being countably complete, must concentrate on one X_m or on the complement of X . Shrinking each X_n again without reducing its cardinality, we arrange that for each fixed n and each fixed $k < l(n)$, as ξ varies over X_n , all the ultrafilters $\mathcal{U}_k(\xi)$ that contain X also contain the same X_m . We write $m(n, k)$ for this m . (If none of these $\mathcal{U}_k(\xi)$ contain X , define $m(n, k) \in \omega - \{n\}$ arbitrarily.) Also, define $S(n) = \{m(n, k) : k < l(n)\}$. Thus, when $\xi \in X_n$, every ultrafilter in $\mathbb{U}_{\nu(\xi)}$ that contains X contains X_m for some $m \in S(n)$. Note that our previous shrinking of the X_n 's ensures that $n \notin S(n)$.

(A technical comment: When we shrink X by shrinking all the X_n 's, the property of an ultrafilter that " $X_m \in \mathcal{U}$ " may be lost, since X_m may shrink to a set not in \mathcal{U} . But, if this happens, then X also shrinks to a set not in \mathcal{U} . Thus, the property "if $X \in \mathcal{U}$ then $X_m \in \mathcal{U}$ " persists under such shrinking. This fact was tacitly used in the shrinking process of the preceding paragraph. It ensures that we can base our decision of how to shrink the X_n 's on our knowledge of which X_m 's are in which ultrafilters, without worrying that the shrinking will alter that knowledge in a way that requires us to revise the shrinking.)

Obtain an infinite subset Y of ω by choosing its elements inductively, in increasing order, so that whenever $n < n'$ are in Y then $n' \notin S(n)$. This is trivial to do, since each $S(n)$ is finite. Shrink X_n to \emptyset for all $n \notin Y$, but leave X_n unchanged for $n \in Y$. Unlike previous shrinkings, this obviously does not maintain $|X_n| = \beta_n$ in general but only for $n \in Y$. That is, however, sufficient to maintain $|X| = \beta$, since Y is cofinal in ω and so the β_n for $n \in Y$ have supremum β . As a result of this last shrinking, we have that, for each $n \in Y$ and each $\xi \in X_n$, each of the ultrafilters $\mathcal{U}_k(\xi) \in \mathbb{U}_{\nu(\xi)}$ that contains X also contains X_m with $m = m(n, k) < n$.

Shrinking the surviving X_n 's further, without reducing their cardinalities,

we can arrange that in formulas (2) and (3) the coefficient $b_{\mathcal{U}_k(\xi)}^{\nu(\xi)}$ depends only on n and k , not on the choice of $\xi \in X_n$. We call this coefficient $b(n, k)$.

We shall now define a certain $x \in \mathbb{Z}^\beta$ with support (the current) X . It will be constant on each X_n with a value z_n to be specified, by induction on n . (Here n ranges over Y , since $X_n = \emptyset$ for $n \notin Y$.) Suppose that integers z_m have already been defined for all $m < n$. Then for $\xi \in X_n$ the sum in formula (3) for $\nu = \nu(\xi)$ is

$$\sum_{\mathcal{U} \in \mathbb{U}_{\nu(\xi)}} b_{\mathcal{U}}^{\nu(\xi)} \cdot \mathcal{U}\text{-lim } x = \sum_{k < l(n)} b(n, k) \cdot \mathcal{U}_k(\xi)\text{-lim } x = \sum_{k < l(n)} b(n, k) \cdot (z_{m(n,k)}|0).$$

Here $(z|0)$ means z or 0 , according to whether $\mathcal{U}_k(\xi)$ contains X (and therefore $X_{m(n,k)}$) or not. So this sum has only finitely many (at most $2^{l(n)}$) possible values. Choose z_n to be an integer greater than the absolute values of these finitely many possible sums. This choice ensures that, in formula (3) for $\nu = \nu(\xi)$ and $\xi \in X_n$, the first term $a_{\xi}^{\nu(\xi)} x(\xi)$ exceeds in absolute value the sum over non-principal ultrafilters. Therefore, $j_{\nu(\xi)}(x) \neq 0$.

But this happens for all $\xi \in X$, so $\text{supp}(j(x))$ has cardinality β , contrary to the fact that $j(x) \in \Pi(\lambda, < \beta)$. This contradiction completes the proof for Case 2.

Case 3: $|X| = \text{cf}(\beta) < \beta$.

We already observed that the basic idea suffices to complete the proof if $|X|$ is smaller than all measurable cardinals. So in the present situation, we may assume that $\text{cf}(\beta)$ is greater than or equal to the first measurable cardinal; in particular it is uncountable.

Let $\mu = \text{cf}(\beta)$ and let $(\beta_i)_{i \in \mu}$ be an increasing μ -sequence of regular, uncountable cardinals with supremum β .

For each $i \in \mu$, there is some $\xi_i \in X$ such that

$$|\{\nu : X \cap F_\nu = \{\xi_i\}\}| \geq \beta_i.$$

Indeed, if there were no such ξ_i , then $\{\nu : |X \cap F_\nu| = 1\}$ would be the union of $|X| = \mu$ sets each of size $< \beta_i$, so it would have cardinality at most $\mu \cdot \beta_i < \beta$, contrary to our original choice of X .

Fix such a ξ_i for each $i \in \mu$. Note that $|\{\nu : X \cap F_\nu = \{\xi_i\}\}|$, though at least β_i by definition, cannot be as large as β , as we remarked when we disposed of the case $|X| = 1$ long ago. So, although the same element can serve as ξ_i for several i 's, it cannot do so for cofinally many $i \in \mu$. So there

are μ distinct ξ_i 's. Passing to a subsequence and re-indexing, we henceforth assume that all the ξ_i are distinct.

Next, fix for each $i \in \mu$ a set $N_i \subseteq \lambda$ of size β_i such that all elements ν of N_i have $X \cap F_\nu = \{\xi_i\}$. Note that the sets N_i are pairwise disjoint.

Shrink X to $\{\xi_i : i \in \mu\}$. This still has cardinality μ and thus has all the properties originally assumed for X .

For each i , shrink N_i , without reducing its cardinality β_i , so that as ν varies over N_i , the cardinality of \mathbb{U}_ν remains constant, say $l(i)$. This shrinking is possible because $\text{cf}(\beta_i) > \omega$. Since μ is uncountable and regular, we can shrink X , without reducing its cardinality, so that $l(i)$ is the same number l for all $\xi_i \in X$. Again, re-index X as $\{\xi_i : i \in \mu\}$ and re-index the β_i and N_i correspondingly. So we can, for each $\nu \in \bigcup_i N_i$, enumerate \mathbb{U}_ν as $\{\mathcal{U}_k(\nu) : k < l\}$.

For each i , choose a uniform ultrafilter \mathcal{V}_i on N_i , and define an ultrafilter \mathcal{W}_i as the limit with respect to \mathcal{V}_i of the ultrafilters $\mathcal{U}_0(\nu)$. That is,

$$A \in \mathcal{W}_i \iff \{\nu : A \in \mathcal{U}_0(\nu)\} \in \mathcal{V}_i.$$

It is well known and easy to check that this \mathcal{W}_i is indeed an ultrafilter. Applying Theorem 3, we obtain $Y \subseteq X$ of cardinality μ , such that for each $\xi_i \in Y$, $Y \notin \mathcal{W}_i$. This means, by definition of \mathcal{W}_i , that we can shrink N_i to a set in \mathcal{V}_i , hence still of size β_i as \mathcal{V}_i is uniform, so that for all ν in the new N_i , $\mathcal{U}_0(\nu)$ doesn't contain Y . Shrink X to Y and reindex as before. We have achieved that, for all i and all $\nu \in N_i$, $X \notin \mathcal{U}_0(\nu)$.

Repeat the process with the subscript 0 of \mathcal{U} replaced in turn by $1, 2, \dots, l-1$. At the end, we have X and N_i 's such that, for all $\xi_i \in X$, all $\nu \in N_i$, and all $\mathcal{U} \in \mathbb{U}_\nu$, $X \notin \mathcal{U}$.

This means that, in formula (3) for $\xi = \xi_i \in X$ and $\nu \in N_i$, if x has support X , then the sum over non-principal ultrafilters vanishes and we reach a contradiction as in the basic idea.

References

- [1] S. Balcerzyk, "On groups of functions defined on Boolean algebras," *Fund. Math.* 50 (1962) 347–367.
- [2] A. Blass and J. Irwin, "Special families of sets and Baer-Specker groups," to appear in *Communications in Algebra*.

- [3] K. Eda, “A Boolean power and a direct product of abelian groups,” *Tsukuba J. Math.* 6 (1982) 187–193.
- [4] L. Fuchs, *Infinite Abelian Groups*, vol. 2, Academic Press (1973).
- [5] T. Jech, *Set Theory*, Academic Press (1978).
- [6] G. Nöbeling, “Verallgemeinerung eines Satzes von Herrn E. Specker,” *Invent. Math.* 6 (1968) 41–55.
- [7] E. Specker, “Additive Gruppen von Folgen ganzer Zahlen,” *Portugaliae Math.* 9 (1950) 131–140.