P-sets and minimal right ideals in $\mathbb{N}^*$

by

W. R. Brian (New Orleans, LA)

Abstract. Recall that a P-set is a closed set X such that the intersection of countably many neighborhoods of X is again a neighborhood of X. We show that if $\mathfrak{t} = \mathfrak{c}$ then there is a minimal right ideal of $(\beta\mathbb{N}, +)$ that is also a P-set. We also show that the existence of such P-sets implies the existence of P-points; in particular, it is consistent with ZFC that no minimal right ideal is a P-set. As an application of these results, we prove that it is both consistent with and independent of ZFC that the shift map and its inverse are (up to isomorphism) the unique chain transitive autohomeomorphisms of $\mathbb{N}^*$.

1. Introduction. The dynamical and algebraic structure of $\beta\mathbb{N}$ and $\mathbb{N}^*$ has played a major part in modern combinatorics and algebra. The minimal right ideals of $\beta\mathbb{N}$—or equivalently, its minimal dynamical subsystems—have a special place in this theory. In this paper we study how these ideals are embedded in $\mathbb{N}^*$.

Our main result is that it is both consistent with and independent of ZFC that some minimal right ideal is also a P-set.

In Section 3, we interpret the minimal right ideals of $\mathbb{N}^*$ as the ultrafilters on a certain partial order. This provides a tool for studying how these ideals are embedded in $\mathbb{N}^*$, and some of the techniques for analyzing the topology of $\mathbb{N}^*$ (the set of all ultrafilters on a different partial order, namely $\mathcal{P}(\omega)/\text{fin}$) carry over. In this way we are able to prove that if $\mathfrak{t} = \mathfrak{c}$ then some minimal right ideal is also a P-set.

In Section 4, we show that if I is a minimal right ideal and a P-set, then there is a finite-to-one $f : \mathbb{N} \to \mathbb{N}$ such that $\beta f(I)$ is a P-point. It follows that it is consistent to have no minimal right ideal be a P-set. We also show in this section that, under ZFC, there is a nowhere dense closed right ideal that is a P-set.

In Section 5, we give an application of these results by obtaining another consistency/independence theorem. The shift map, the canonical dynamical structure on $\mathbb{N}^*$ obtained by extending the map $n \mapsto n + 1$ to $\beta\mathbb{N}$, is shown to be a chain transitive map. We will prove that it is consistent with and independent of ZFC that the shift map and its inverse are the only chain transitive autohomeomorphisms of $\mathbb{N}^*$. The proof of independence will use the fact that some nowhere dense right ideal is a P-set, together with Parovichenko's topological characterization of $\mathbb{N}^*$ under CH.

2. Notation and preliminaries

Ultrafilters. Let $(\mathbb{P}, \leq)$ be a partial order. A *filter* on $\mathbb{P}$ is a subset $\mathcal{F}$ of $\mathbb{P}$ satisfying

- (nontriviality) $\emptyset \neq \mathcal{F} \neq \mathbb{P}$;
- (upward heredity) if $a \in \mathcal{F}$ and $a \leq b$ then $b \in \mathcal{F}$;
- (finite intersection property) if $a, b \in \mathcal{F}$ then there is some $c \in \mathcal{F}$ with $c \leq a$ and $c \leq b$.

An *ultrafilter* on $\mathbb{P}$ is a maximal filter: a filter $\mathcal{F}$ such that no set properly containing $\mathcal{F}$ is also a filter. Ultrafilters on the partial order $(\mathcal{P}(\mathbb{N}), \subseteq)$ are perhaps most familiar, and here we will call these simply *ultrafilters on* $\mathbb{N}$ or, context permitting, ultrafilters. A *free* (ultra)filter on $\mathbb{N}$ is one that extends the filter of cofinite sets (or equivalently, a filter with empty intersection).

If $\mathcal{F}$ satisfies the nontriviality and finite intersection properties, then $\mathcal{F}$ is called a *filter base*, and $\{b \in P : a \leq b \text{ for some } a \in \mathcal{F}\}$ is a filter, namely the filter *generated by* $\mathcal{F}$. A *chain* in $\mathbb{P}$ is a totally ordered subset of $\mathbb{P}$. Note that every chain is a filter base.

If $\mathcal{F}$ satisfies the nontriviality property and is upwards hereditary, then $\mathcal{F}$ is an *up-set* of $\mathbb{P}$ (the up-sets of $(\mathcal{P}(\mathbb{N}), \subseteq)$ are often referred to as *Furstenberg families* in the dynamical systems literature; the up-sets of $\mathcal{P}(\omega)/\text{fin}$ are often referred to as *semifilters* in the set theory literature).

Following convention, if $A, B \subseteq \mathbb{N}$ we write $A \subseteq^* B$ whenever $A \setminus B$ is finite and we write $A =^* B$ whenever $A \subseteq^* B$ and $B \subseteq^* A$; $\mathcal{P}(\omega)/\text{fin}$ is the set of $=^*$-equivalence classes of infinite subsets of $\mathbb{N}$, naturally ordered by $\subseteq^*$. It is a well-studied Boolean algebra, and the ultrafilters on the partial order $\mathcal{P}(\omega)/\text{fin} - \{0\}$ comprise the well-known space $\mathbb{N}^*$. From now on, we will abuse notation in the standard way and write $\mathcal{P}(\omega)/\text{fin}$ to mean the partial order $\mathcal{P}(\omega)/\text{fin} - \{0\}$. We will also sometimes treat subsets of $\mathbb{N}$ as elements of $\mathcal{P}(\omega)/\text{fin}$, even though the elements of $\mathcal{P}(\omega)/\text{fin}$ are really equivalence classes.

Given $(\mathbb{P}, \leq)$ and $A \subseteq \mathbb{P}$, $p \in \mathbb{P}$ is a (*lower*) *bound* for A provided $p \leq a$ for every $a \in A$; A is *unbounded* if it has no bound. Note: "pseudo-intersection" is often used synonymously with "lower bound" when working with $\mathcal{P}(\omega)/\text{fin}$.

Some small cardinals. Much has been written about the "small cardinal" numbers, of which there are many. Here we define four, three old and one new (the introduction of a new one can be forgiven, since we will show in Theorem 3.4 it is really just $\mathfrak{t}$ in disguise). Our definitions are not quite standard. This is intentional, to emphasize that our main theorem is really about a certain partial order Θ. In fact, the proof of our main theorem can be viewed as the translation to Θ of a well-known proof done in $\mathcal{P}(\omega)/\mathrm{fin}$.

Recall that $A \subseteq \mathbb{N}$ is *thick* if it contains arbitrarily long intervals. Equivalently, A fails to be thick if and only if there is some k such that $\bigcap_{n \leq k}(A - n)$ is empty. The thick sets form an up-set of $\mathcal{P}(\mathbb{N})$. If $A =^* B$, then A is thick if and only if B is thick, so the (equivalence classes of) thick sets also form a natural up-set of $\mathcal{P}(\omega)/\mathrm{fin}$. Let Θ denote the family of thick sets modulo $=^*$, partially ordered by $\subseteq^*$. As with $\mathcal{P}(\omega)/\mathrm{fin}$, we will sometimes treat subsets of $\mathbb{N}$ as members of Θ. We refer to the (ultra)filters on $(\Theta, \subseteq^*)$ as Θ-*(ultra)filters*.

Recall that $\mathbb{N}^\mathbb{N}$ denotes the set of all functions $\mathbb{N} \to \mathbb{N}$, and if $f, g \in \mathbb{N}^\mathbb{N}$, then $f <^* g$ if and only if $\{n \in \mathbb{N} : f(n) \not< g(n)\}$ is finite.

We now define the small cardinals we will need:

- $\mathfrak{t}_\Theta$ is the smallest size of an unbounded chain in Θ.
- $\mathfrak{t}$ is the smallest size of an unbounded chain in $\mathcal{P}(\omega)/\mathrm{fin}$.
- $\mathfrak{p}$ is the smallest size of an unbounded filter base in $\mathcal{P}(\omega)/\mathrm{fin}$.
- $\mathfrak{b}$ is the smallest size of an unbounded subset of $(\mathbb{N}^\mathbb{N}, >^*)$.

This way, it is easy to see that $\mathfrak{t}_\Theta$ is simply the analogue of $\mathfrak{t}$ in Θ. While the definition of $\mathfrak{t}_\Theta$ is ours, the last three cardinals here are standard, and we leave it to the reader to verify that these definitions are equivalent to the usual ones given in [vD] or [Bl2].

Observe that the partial order $(\mathbb{N}^\mathbb{N}, >^*)$ has the property of being a filter. Using this, $\mathfrak{b}$ can also be characterized as the smallest size of an unbounded chain in $(\mathbb{N}^\mathbb{N}, >^*)$ or an unbounded filter base in $(\mathbb{N}^\mathbb{N}, >^*)$ (see also [vD, Theorem 3.3]). Thus $\mathfrak{b}$ is also just the analogue of $\mathfrak{t}$ (or of $\mathfrak{p}$) in a different partial order.

$\beta\mathbb{N}$ **and Stone duality.** The set of all ultrafilters on $\mathbb{N}$ is denoted $\beta\mathbb{N}$. As usual, we identify the principal ultrafilter $\{A \subseteq \mathbb{N} : n \in A\}$ with n, so that $\mathbb{N} \subseteq \beta\mathbb{N}$. For each $A \subseteq \mathbb{N}$, $\overline{A} = \{p \in \beta\mathbb{N} : A \in p\}$, and taking the sets of this form as the basis for a topology, $\beta\mathbb{N}$ is the Stone–Čech compactification of $\mathbb{N}$. This is the unique compactification of $\mathbb{N}$ with the extension property: for any compact Hausdorff space K and any map $f : \mathbb{N} \to K$, there is a (unique) continuous extension $\beta f : \beta\mathbb{N} \to K$. We write $\mathbb{N}^* = \beta\mathbb{N} \setminus \mathbb{N}$ for the set of all free ultrafilters on $\mathbb{N}$, and, for every $A \subseteq \mathbb{N}$, $A^* = \overline{A} \cap \mathbb{N}^*$.

If $\mathcal{F}$ is a free (ultra)filter on $\mathbb{N}$, then $\mathcal{F}$ restricts to an (ultra)filter on $\mathcal{P}(\omega)/\mathrm{fin}$. Conversely, an (ultra)filter on $\mathcal{P}(\omega)/\mathrm{fin}$ gives rise to a unique free (ultra)filter. In what follows we will freely conflate these two notions. Topologically this makes no difference, as $A =^* B$ if and only if $A^* = B^*$.

If $\mathcal{F}$ is a free filter, then $\mathcal{F}^* = \bigcap_{A \in \mathcal{F}} A^*$ is a nonempty closed subset of $\mathbb{N}^*$. Conversely, for every closed $X \subseteq \mathbb{N}^*$ there is a unique free filter $\mathcal{F}$ such that $\mathcal{F}^* = X$, namely $\mathcal{F} = \{A \subseteq \mathbb{N} : X \subseteq A^*\}$. This correspondence, part of what is called Stone duality, will be used frequently (often implicitly) in what follows.

A closed subset X of $\mathbb{N}^*$ is a *P-set* if, for any countable collection $\mathcal{A}$ of open sets containing X, X is in the interior of $\bigcap \mathcal{A}$. If $X = \{x\}$ is a P-set, then x is called a *P-point*. If $\mathcal{F}$ is a filter such that $\mathcal{F}^*$ is a P-set, we say that $\mathcal{F}$ is a *P-filter*. By Stone duality, $\mathcal{F}$ is a P-filter if and only if every countable subset of $\mathcal{F}$ has a lower bound in $\mathcal{F}$.

Dynamics and algebra. A *dynamical system* is a pair (X, f), where X is a compact Hausdorff space and $f : X \to X$ is continuous. For our purposes, X will usually be $\beta\mathbb{N}$ or $\mathbb{N}^*$.

For each $p \in \mathbb{N}^*$, define $\sigma(p)$ to be the unique ultrafilter generated by $\{A + 1 : A \in p\}$. This is the *shift map* on $\beta\mathbb{N}$, and whenever we speak of $\beta\mathbb{N}$ as a dynamical system it is understood that we are talking about the shift map. The shift map is the unique continuous extension to $\beta\mathbb{N}$ of the map on $\mathbb{N}$ given by $n \mapsto n + 1$.

If $\langle x_n : n \in \mathbb{N} \rangle$ is a countable sequence of points in any compact Hausdorff space and if $p \in \beta\mathbb{N}$, we define $p\text{-}\lim_{n \in \mathbb{N}} x_n$ to be the unique element of $\bigcap_{A \in p} \overline{\{x_n : n \in A\}}$. Then, for any $p, q \in \beta\mathbb{N}$, define $p + q = q\text{-}\lim_{n \in \mathbb{N}} \sigma^n(p)$.

Some authors (see, e.g., [Be]), define $+$ differently on $\beta\mathbb{N}$:

$$p + q = \big\{ A \subseteq \mathbb{N} : \{n \in \mathbb{N} : (A - n) \in p\} \in q \big\}.$$

This definition is equivalent to the one we give here (the equivalence is well known; see [Bl1] for some discussion). Warning: some authors, including some that we cite below, exchange the roles of p and q in this definition.

The map $q \mapsto p + q$ (with p fixed) is continuous, but the map $p \mapsto p + q$ (with q fixed) is not. This makes $\beta\mathbb{N}$ a left-topological semigroup. Its *right ideals* are those $I \subseteq \beta\mathbb{N}$ such that $I + \beta\mathbb{N} \subseteq I$. A *minimal* right ideal is a right ideal that does not properly contain any other right ideal.

For any dynamical system (X, f), Y is a *subsystem* of X provided that Y is closed and $f(Y) \subseteq Y$. A subsystem is *minimal* if it does not properly contain any other subsystems.

A simple application of Zorn's Lemma shows that every subsystem of $\beta\mathbb{N}$ contains a minimal subsystem, and every right ideal contains a minimal right ideal. In $\beta\mathbb{N}$, these are related by the following:

PROPOSITION 2.1. *X is a subsystem of* $(\beta\mathbb{N}, \sigma)$ *if and only if it is a closed right ideal of* $(\beta\mathbb{N}, +)$; *and X is a minimal subsystem if and only if X is a minimal right ideal.*

Proof. See, e.g., [Be, Theorem 2.1]. ∎

3. Θ-ultrafilters. The following two lemmas make the connection between Θ-ultrafilters and minimal right ideals. These lemmas are hinted at in Theorem 2.9(c) of [BHM], which states that if $A \subseteq \mathbb{N}$ then A^* contains a right ideal of $\mathbb{N}^*$ if and only if A is thick.

LEMMA 3.1. *Let $\mathcal{F}$ be a free filter. Then $\sigma(\mathcal{F}^*) \subseteq \mathcal{F}^*$ if and only if for every $A \in \mathcal{F}$, $A - 1 \in \mathcal{F}$. Further, $\mathcal{F}^* = \sigma(\mathcal{F}^*)$ if and only if, in addition, $A + 1 \in \mathcal{F}$ for every $A \in \mathcal{F}$.*

Proof. The "only if" directions of both claims follow easily from Stone duality and the continuity of σ^{-1} and σ:

$$\sigma^{-1}(\mathcal{F}^*) = \sigma^{-1}\Big(\bigcap_{A \in \mathcal{F}} A^*\Big) = \bigcap_{A \in \mathcal{F}} \sigma^{-1}(A^*) = \bigcap_{A \in \mathcal{F}} (A - 1)^*,$$

$$\sigma(\mathcal{F}^*) = \sigma\Big(\bigcap_{A \in \mathcal{F}} A^*\Big) = \bigcap_{A \in \mathcal{F}} \sigma(A^*) = \bigcap_{A \in \mathcal{F}} (A + 1)^*.$$

We can then use these equalities to prove the "if" direction of these claims: if $A - 1 \in \mathcal{F}$ for every $A \in \mathcal{F}$ then

$$\sigma^{-1}(\mathcal{F}^*) = \bigcap_{A \in \mathcal{F}} (A - 1)^* \supseteq \bigcap_{A \in \mathcal{F}} A^* = \mathcal{F}^*,$$

which implies $\sigma(\mathcal{F}^*) \subseteq \mathcal{F}^*$. If also $A - 1 \in \mathcal{F}$ for every $A \in \mathcal{F}$ then

$$\sigma(\mathcal{F}^*) = \bigcap_{A \in \mathcal{F}} (A + 1)^* \supseteq \bigcap_{A \in \mathcal{F}} A^* = \mathcal{F}^*.$$

In the latter case, we clearly have $\mathcal{F}^* = \sigma(\mathcal{F}^*)$. ∎

LEMMA 3.2. *Let $\mathcal{F}$ be a free filter. Then $\mathcal{F}^*$ is a minimal right ideal if and only if $\mathcal{F}$ is a Θ-ultrafilter.*

Proof. First we recall a well-known fact: every minimal right ideal in $\mathbb{N}^*$ is closed. To see this, let I be a minimal right ideal and let $p \in I$. Then $p + \mathbb{N}^* \subseteq I$ is a right ideal, so we must have $p + \mathbb{N}^* = I$. Since the map $q \mapsto p + q$ is continuous and $\mathbb{N}^*$ is compact, $p + \mathbb{N}^*$ is compact, hence closed.

If $\mathcal{F}^*$ is a closed right ideal, it follows directly from [BHM, Theorem 2.9(c)] that $\mathcal{F} \subseteq \Theta$, i.e., $\mathcal{F}$ is a Θ-filter. By Proposition 2.1, the closed right ideals of $\mathbb{N}^*$ are precisely those sets $\mathcal{F}^*$ such that $\mathcal{F}$ is a Θ-filter and $\sigma(\mathcal{F}^*) \subseteq \mathcal{F}^*$. Thus the minimal (closed) right ideals are the sets $\mathcal{F}^*$ where $\mathcal{F}$ is maximal among all Θ-filters satisfying $\sigma(\mathcal{F}^*) \subseteq \mathcal{F}^*$.

To prove the lemma, we will show that every Θ-ultrafilter $\mathcal{F}$ satisfies $\sigma(\mathcal{F}^*) \subseteq \mathcal{F}^*$. Since the Θ-ultrafilters are by definition just the maximal Θ-filters, this shows that they are also, *ipso facto*, the maximal Θ-filters satisfying $\sigma(\mathcal{F}^*) \subseteq \mathcal{F}^*$.

Let $\mathcal{F}$ be a Θ-ultrafilter and $A, B \in \mathcal{F}$. Then $A \cap B$ is thick, and it is clear that if there is an interval of length k in $A \cap B$ then there is an interval of length $k - 1$ in $(A - 1) \cap B$. Thus $(A - 1) \cap B$ is also thick. Since this is true for any $B \in \mathcal{F}$, it follows from the maximality of Θ-ultrafilters that $A - 1 \in \mathcal{F}$. By Lemma 3.1, $\sigma(\mathcal{F}^*) \subseteq \mathcal{F}^*$. ∎

We thank the referee for pointing out that Lemma 3.2 has a natural extension to arbitrary semigroups. We have not defined all of the concepts involved in this extension (see [HS] instead), but we will state it here for the interested reader:

PROPOSITION 3.3. *Let $(S, \cdot)$ be an arbitrary semigroup and let βS have the operation making $(\beta S, \cdot)$ a left topological semigroup with S contained in its topological center. Let R be a closed subset of βS and let $\mathcal{F} = \{A \subseteq S : R \subseteq \overline{A}\}$. Then R is a minimal right ideal of βS if and only if $\mathcal{F}$ is maximal in $\{\mathcal{G} : \mathcal{G}$ is a filter on S and $(\forall A \in \mathcal{G})(A$ is thick$)\}$.*

By Lemma 3.2, some minimal right ideal is a P-set if and only if some Θ-ultrafilter is also a P-filter. By Stone duality, a P-point exists if and only if some ultrafilter on $\mathcal{P}(\omega)/\text{fin}$ is also a P-filter. In either case we are just looking for P-ultrafilters on a given partial order. This is what we mean when we say that our main result is a translation to Θ of a theorem in $\mathcal{P}(\omega)/\text{fin}$.

The first proof of the existence of P-points under MA, given by Booth [Bth] as a modification of Rudin's [R] proof from CH, uses $\mathfrak{t} = \mathfrak{c}$ to build a chain in $\mathcal{P}(\omega)/\text{fin}$ such that the filter generated by this chain is both an ultrafilter and a P-filter.

To prove that it is possible to have P-sets that are minimal right ideals, we will do the same thing as Booth, only replacing $\mathcal{P}(\omega)/\text{fin}$ with Θ and $\mathfrak{t}$ with $\mathfrak{t}_\Theta$. The only real difference is that, instead of showing MA implies $\mathfrak{t}_\Theta = \mathfrak{c}$ directly, we will prove that $\mathfrak{t}_\Theta = \mathfrak{t}$.

THEOREM 3.4. $\mathfrak{t}_\Theta = \mathfrak{t}$.

Proof. We will prove that $\mathfrak{p} \leq \mathfrak{t}_\Theta \leq \mathfrak{t}$. This is enough because, by a deep new result of Malliaris and Shelah, $\mathfrak{p} = \mathfrak{t}$ (see [MS]).

To prove $\mathfrak{t}_\Theta \leq \mathfrak{t}$ we make use of the standard trick: supposing $\kappa < \mathfrak{t}_\Theta$, we will show $\kappa < \mathfrak{t}$. Suppose $\kappa < \mathfrak{t}_\Theta$ and let $T = \{A_\alpha : \alpha < \kappa\}$ be a chain in $\mathcal{P}(\omega)/\text{fin}$. For each n, let $I_n = [2^n, 2^{n+1} - 1]$. For each $\alpha < \kappa$, let $B_\alpha = \bigcup_{n \in A} I_n$. Clearly each B_α is thick and $\{B_\alpha : \alpha < \kappa\}$ is a chain in Θ. Since $\kappa < \mathfrak{t}_\Theta$, there is some B_κ such that $B_\kappa \subseteq^* B_\alpha$ for all $\alpha < \kappa$. Setting

$A_\kappa = \{n : B_\kappa \cap I_n \neq \emptyset\}$, we easily check that $A_\kappa \subseteq^* A_\alpha$ for all $\alpha < \kappa$. Since T was arbitrary, $\kappa < \mathfrak{t}$.

We use the same trick to prove $\mathfrak{p} \leq \mathfrak{t}_\Theta$. Suppose $\kappa < \mathfrak{p}$ and let $\{T_\alpha : \alpha < \kappa\}$ be a chain in Θ; we will show that there is some $T_\kappa \in \Theta$ such that $T_\kappa \subseteq^* T_\alpha$ for every α.

Let $\mathcal{T} = \{T_\alpha - n : \alpha < \kappa,\, n \in \mathbb{N}\}$. We claim that the closure of $\mathcal{T}$ under finite intersections is a filter base. Fix $\alpha_0, \ldots, \alpha_m$ and $n_0, \ldots, n_m$. There is a thick set $T \subseteq \bigcap_{k \leq m} T_{\alpha_k}$; in fact, some cofinite subset of $T_{\max\{\alpha_0, \ldots, \alpha_n\}}$ will do. Let $N = \max\{n_0, \ldots, n_m\}$. Because T is thick, there is an infinite $T' \subseteq T$ such that for every $\ell \in T'$ we have $[\ell, \ell + N] \subseteq T$. But if $[\ell, \ell + N] \subseteq T$ then $\ell \in \bigcap_{k \leq m}(T_{\alpha_k} - n_k)$; thus $\bigcap_{k \leq m}(T_{\alpha_k} - n_k)$ is infinite.

The closure of $\mathcal{T}$ under finite intersections is a filter base, and it has size $|\mathcal{T}| = \kappa$. Since $\kappa < \mathfrak{p}$, this filter base has a lower bound A in $\mathcal{P}(\omega)/\text{fin}$. Let $\{a_n : n \in \mathbb{N}\}$ be an increasing enumeration of A.

For each $\alpha < \kappa$ and $m \in \mathbb{N}$, there is some minimal $\ell \in \mathbb{N}$ such that

$$\{a_n : n \geq \ell\} \subseteq \bigcap_{k \leq m}(T_\alpha - k).$$

Let $f_\alpha(m) = \ell$. Equivalently, $f_\alpha(m)$ is the least element of $\mathbb{N}$ such that if $n \geq f_\alpha(m)$ then $[a_n, a_n + m] \subseteq T_\alpha$. It is well known that $\mathfrak{p} \leq \mathfrak{b}$ (see, e.g., [vD]). Therefore $\kappa < \mathfrak{b}$ and there is some $f : \mathbb{N} \to \mathbb{N}$ such that, for all $\alpha < \kappa$, $f_\alpha <^* f$. Let

$$T_\kappa = \bigcup_{n \in \mathbb{N}} [a_{f(n)}, a_{f(n)} + n].$$

Clearly $T_\kappa \in \Theta$. Fix $\alpha < \kappa$. There is some $N \in \mathbb{N}$ such that $f_\alpha(n) < f(n)$ for $n \geq N$. But if $f_\alpha(n) < f(n)$ then $[a_{f(n)}, a_{f(n)} + n] \subseteq T_\alpha$. Thus $T_\kappa \subseteq^* T_\alpha$, as required. ∎

We leave it an open question whether there is an "elementary" proof that $\mathfrak{t}_\Theta = \mathfrak{t}$, i.e., one that avoids the power of Malliaris and Shelah's equality. We also leave as an exercise the (easy) proof that MA can be used directly to show $\mathfrak{t}_\Theta = \mathfrak{c}$.

THEOREM 3.5. *If* $\mathfrak{t} = \mathfrak{c}$ *then there is a minimal right ideal in* $\mathbb{N}^*$ *that is also a P-set.*

Proof. Using Lemma 3.2, it suffices to construct a Θ-ultrafilter that is also a P-filter.

Fix an enumeration $\{A_\alpha : \alpha < \mathfrak{c}\}$ of Θ. We construct a (reverse well-ordered) chain $T = \{X_\alpha : \alpha < \mathfrak{c}\}$ in Θ as follows. Set $X_0 = \mathbb{N}$. If X_α has already been defined, let $X_{\alpha+1} = X_\alpha \cap A_\alpha$ if $X_\alpha \cap A_\alpha \in \Theta$, and otherwise let $X_{\alpha+1} = X_\alpha$. For limit α, let X_α be any lower bound in $(\Theta, \subseteq^*)$ of $\{X_\beta : \beta < \alpha\}$; a lower bound exists because $\alpha < \mathfrak{t}_\Theta$.

The chain T constructed in the previous paragraph is a filter base; let $\mathcal{F}$ be the filter generated by T. If $A_\alpha \in \Theta$, then either $A_\alpha \subseteq X_{\alpha+1}$, which implies $A_\alpha \in \mathcal{F}$, or $A_\alpha \cap X_\alpha$ is not thick. Thus $\mathcal{F}$ is a Θ-ultrafilter. If $\mathcal{A}$ is a countable subset of $\mathcal{F}$ then, shrinking the elements of $\mathcal{A}$ if necessary, we may assume $\mathcal{A} \subseteq T$ and write $\mathcal{A} = \{X_{\alpha_n} : n \in \omega\}$. Since the cofinality of $\mathfrak{c}$ is uncountable, there is some $\alpha < \mathfrak{c}$ with $\alpha_n < \alpha$ for all n. Then X_α is a pseudo-intersection for $\mathcal{A}$ in $\mathcal{F}$, so that $\mathcal{F}$ is a P-filter. ∎

REMARK 3.6. The minimal right ideal constructed in the previous proof is actually a $P_{\mathfrak{c}}$-set (the proof implies automatically that it is a $P_{\mathrm{cf}(\mathfrak{c})}$-set, and $\mathfrak{t}$ is a regular cardinal).

As remarked in the final section of [Bl1], the elements of minimal right ideals are never P-points. In fact, if I is a minimal right ideal and $p \in I$, then $\{\sigma^n(p) : n \in \mathbb{N}\}$ is dense in I by Proposition 2.1; since the minimal right ideals are separable with no isolated points, no point of one is even a weak P-point. Our result states that (sometimes, in some models) this is the only reason a minimal ultrafilter fails to be a P-point.

4. More on P-set right ideals. We begin this section by proving that the conclusion of Theorem 3.5 is consistently false.

LEMMA 4.1. *If $f : \mathbb{N} \to \mathbb{N}$ is finite-to-one, then βf maps closed P-sets to closed P-sets.*

Proof. This lemma is a well-known bit of folklore, but we will give the short proof here for completeness.

Let $\mathcal{F}$ be a filter with $\mathcal{F}^*$ a closed P-set, and let $\mathcal{G}$ denote the unique filter such that $\beta f(\mathcal{F}^*) = \mathcal{G}^*$. One can easily check that

$$\mathcal{G} = \{A \subseteq \mathbb{N} : f^{-1}(A) \in \mathcal{F}\}.$$

Let $\mathcal{A} \subseteq \mathcal{G}$ be countable. Since $\mathcal{F}$ is a P-filter, $\mathcal{F}$ contains a pseudo-intersection B for $\{f^{-1}(A) : A \in \mathcal{A}\}$. Now $f(B) = \{n : B \cap f^{-1}(n) \neq \emptyset\}$ is a pseudo-intersection for $\mathcal{A}$, so $\mathcal{G}$ is a P-filter as well. ∎

LEMMA 4.2. *There is a finite-to-one map $f : \mathbb{N} \to \mathbb{N}$ such that if I is a minimal right ideal then $\beta f(I)$ is a single point. Furthermore, for every $p \in \mathbb{N}^*$ there is some minimal right ideal I such that $\beta f(I) = p$.*

Proof. Let $\{a_n : n \in \mathbb{N}\}$ be an increasing sequence of natural numbers such that $\lim_{n\to\infty}(a_{n+1} - a_n) = \infty$, and let f be the map that sends m to the greatest n such that $a_n \leq m$. We will show that f has the required property.

For each n, let $I_n = f^{-1}(n)$. Let $\mathcal{F}$ be a Θ-ultrafilter. Letting

$$\beta f(\mathcal{F}) = \{A \subseteq \mathbb{N} : f^{-1}(A) \in \mathcal{F}\}$$

as in the proof of Lemma 4.1, we see it suffices to check that $\beta f(\mathcal{F})$ is an ultrafilter on $\mathbb{N}$. We will show that for every $A \subseteq \mathbb{N}$ either $f^{-1}(A) \in \mathcal{F}$ or $f^{-1}(\mathbb{N} \setminus A) \in \mathcal{F}$.

Suppose this is not the case, and let A be such that neither $f^{-1}(A)$ nor $f^{-1}(\mathbb{N} \setminus A)$ is in $\mathcal{F}$. As $\mathcal{F}$ is a Θ-ultrafilter, this means that there are $B, C \in \mathcal{F}$ such that neither $f^{-1}(A) \cap B$ nor $f^{-1}(\mathbb{N} \setminus A) \cap C$ is thick. $\mathcal{F}$ is a filter, so $B \cap C \in \mathcal{F}$, and neither $D = B \cap C \cap f^{-1}(A)$ nor $E = B \cap C \cap f^{-1}(\mathbb{N} \setminus A)$ is thick. Fix k such that neither D nor E contains an interval of length $\geq k$. There is some M such that if $n > M$ then $a_{n+1} - a_n > k$. If I is an interval in $D \cup E$ with $\min I \geq a_M$, then I has length at most $2k$ (otherwise either D or E contains an interval of length greater than k). Thus $D \cup E$ is not thick. Since

$$D \cup E = (B \cap C \cap f^{-1}(A)) \cup (B \cap C \cap f^{-1}(\mathbb{N} \setminus A)) = B \cap C$$

and $B \cap C \in \mathcal{F}$, this is a contradiction.

For the second claim, observe that for any $p \in \mathbb{N}^*$,

$$\left\{ \bigcup_{n \in A} [a_n, a_{n+1} - 1] : A \in p \right\}$$

is a Θ-filter. If $\mathcal{F}$ is a Θ-ultrafilter extending this filter, then $\mathcal{F}^*$ is a minimal right ideal with $\beta f(\mathcal{F}^*) = p$. ∎

We remark that Blass and Hindman in [BlH] have proved a similar result. They show that if MA holds, then the (finite-to-one) map f on $\mathbb{N}$ that sends n to the leftmost 1 in its binary expansion has the property that βf sends some idempotent ultrafilter to a Ramsey ultrafilter. Our lemma shows that if p is any ultrafilter, then this map (or one of many others like it) will send some minimal idempotent to p.

THEOREM 4.3. *It is consistent with ZFC that no minimal right ideal is a P-set.*

Proof. By Lemmas 4.1 and 4.2, if some minimal right ideal is a P-set then there is a P-point in $\mathbb{N}^*$. By a famous result of Shelah, it is consistent with ZFC that $\mathbb{N}^*$ has no P-points (see [W] or [S2]). ∎

The proof of Theorem 4.3 suggests the following problem:

PROBLEM 4.4. *Is it consistent that* $\mathbb{N}^*$ *has a P-point but that no minimal right ideal is a P-set?*

If $X, Y \subseteq Z$, we say that X and Y are *homeomorphically embedded* in Z if there is a homeomorphism $Z \to Z$ that restricts to a homeomorphism $X \to Y$. Thus, for example, Z is homogeneous if and only if all of its points are homeomorphically embedded.

Recall that p is *idempotent* if and only if $p + p = p$. The idempotents of $\mathbb{N}^*$ admit a natural partial order as follows: if p and q are idempotent ultrafilters, then $p \leq q$ if and only if $p + q = q + p = p$. It is known (see, e.g., [Be, Theorem 2.7]) that an idempotent p is minimal with respect to this order if and only if p belongs to some minimal right ideal. Such ultrafilters are called *minimal idempotents*. It is also known that if q is any idempotent then there is a minimal idempotent p with $p \leq q$.

We now give three corollaries of our Theorem 3.5:

COROLLARY 4.5. *Suppose* $\mathfrak{t} = \mathfrak{c}$. *Then:*

(1) *The minimal right ideals are not homeomorphically embedded in* $\mathbb{N}^*$.
(2) *There is a minimal right ideal* $I \subseteq \beta\mathbb{N}$ *such that, for any* $p, q \in \beta\mathbb{N}$, $p + q \in I$ *if and only if* $p \in I$.
(3) *There is an idempotent ultrafilter that is both minimal and maximal with respect to* $\leq$.

Proof. (1) It suffices to show that some minimal right ideal is not a P-set. Let $K = \bigcup\{I \subseteq \mathbb{N}^* : I$ is a minimal right ideal$\}$ and fix $p \in \mathbb{N}^* \setminus K$. Then $p + \beta\mathbb{N} = \overline{\{\sigma^n(p) : n \in \mathbb{N}\}}$ is clearly a right ideal and, by Zorn's Lemma, must contain a minimal right ideal I. Since I is closed under σ^{-1} (by Proposition 2.1), and $p \notin I$, we have $\sigma^n(p) \notin I$ for all n. It follows that I is not a P-set.

(2) Let I be a P-set and a minimal right ideal. If $p \in I$, then $p + q \in I$ because I is a right ideal. If $p \notin I$, then $p + q = q\text{-}\lim_{n \in \mathbb{N}} \sigma^n(p)$ is an element of $\overline{\{\sigma^n(p) : n \in \mathbb{N}\}}$. Since, as in (1), $\sigma^n(p) \notin I$ for all n, we have $I \cap \overline{\{\sigma^n(p) : n \in \mathbb{N}\}} = \emptyset$. In particular, $p + q \notin I$.

(3) Let I be a minimal right ideal satisfying (2), and let $p \in I$ be idempotent. Let q be any idempotent other than p. If $q \in I$ then q is $\leq$-minimal, so $p \not\leq q$. If $q \notin I$ then $q + p \notin I$ by (2), in which case $p \not\leq q$. ∎

We remark that (2) and (3) are already known to follow from $\mathfrak{t} = \mathfrak{c}$ (see [Z]), though our proof is different. Also, the consistency of (1) is easily inferred from known results: there are $2^{\mathfrak{c}}$ minimal right ideals and it is consistent to have only $\mathfrak{c}$ autohomeomorphisms of $\mathbb{N}^*$ (in fact, this follows from PFA; see [SS]). That (1) follows from $\mathfrak{t} = \mathfrak{c}$ is a new result, and we do not know whether it follows from ZFC.

PROBLEM 4.6. *Is any of the three propositions listed in Corollary* 4.5 *provable from ZFC?*

It may be possible to solve all three parts of this question at once. Recall that a *weak P-set* is a closed $X \subseteq \mathbb{N}^*$ such that if $Y \subseteq \mathbb{N}^* \setminus X$ is countable then $X \cap \overline{Y} = \emptyset$. The proof of Corollary 4.5 does not require that we have a minimal right ideal that is a P-set: having a minimal right ideal that is a weak P-set is enough. In ZFC, it is impossible to prove that P-points exist

(by the aforementioned result of Shelah), but it is possible to prove that weak P-points exist (a result of Kunen [K]). Perhaps a similar situation holds for the minimal right ideals.

PROBLEM 4.7. *Is some minimal right ideal a weak P-set?*

While ZFC is not strong enough to prove that some minimal right ideal is a P-set, we can show that some "small" right ideal is a P-set.

THEOREM 4.8. *There is a nonempty, nowhere dense P-set $X \subseteq \mathbb{N}^*$ such that $\sigma(X) = X$. In particular, there is a nowhere dense P-set that is also a closed right ideal.*

Proof. The second claim of the theorem follows from the first claim and Proposition 2.1.

By Lemma 3.1, it suffices to find a Θ-filter $\mathcal{F}$ such that $\mathcal{F}^*$ is nowhere dense, every countable subset of $\mathcal{F}$ has a pseudo-intersection in $\mathcal{F}$, and, for every $A \in \mathcal{F}$, we have $A + 1, A - 1 \in \mathcal{F}$.

Let $B = \{b_n : n \in \mathbb{N}\}$ satisfy $\lim_{n\to\infty}(b_{n+1} - b_n) = \infty$. Let $\mathcal{G}$ be a filter such that $\mathcal{G}^*$ is a nowhere dense P-set (it is easy to prove that such a $\mathcal{G}$ exists: take a maximal chain in $\mathcal{P}(\omega)/\text{fin}$ and let $\mathcal{G}$ be the filter generated by it). Let $\mathcal{B}$ denote the family of all functions $f : \mathbb{N} \to \mathbb{N}$ such that $\lim_{n\to\infty} f(n) = \infty$. By a standard diagonalization argument, if $f_0, f_1, f_2, \ldots \in \mathcal{B}$ then there is some $f \in \mathcal{B}$ such that $f <^* f_n$ for every n.

For every $G \in \mathcal{G}$ and $f \in \mathcal{B}$, let

$$X(G, f) = \bigcup_{n \in G} [b_n - f(n), b_n + f(n)].$$

Let $G_0, G_1 \in \mathcal{G}$ and $f_0, f_1 \in \mathcal{B}$. Then $G_2 = G_0 \cap G_1 \in \mathcal{G}$, and if we define $f_2 : \mathbb{N} \to \mathbb{N}$ by $f_2(n) = \min\{f_0(n), f_1(n)\}$ then $f_2 \in \mathcal{B}$. Clearly $X(G_2, f_2) \subseteq X(G_0, f_0) \cap X(G_1, f_1)$. Thus $\{X(G, f) : G \in \mathcal{G}, f \in \mathcal{B}\}$ is a filter base; let $\mathcal{F}$ denote the filter generated by it. Clearly, $\mathcal{F}$ is a Θ-filter; we now show that it has the other required properties.

For $f \in \mathcal{B}$, let $f^-(n) = f(n) - 1$ and note that $f^- \in \mathcal{B}$ whenever $f \in \mathcal{B}$. If $G \in \mathcal{G}$ and $f \in \mathcal{B}$, it follows from our definitions that $X(G, f) + 1 \supseteq X(G, f^-)$ and $X(G, f) - 1 \supseteq X(G, f^-)$. Consequently, for every $A \in \mathcal{F}$ we have $A + 1, A - 1 \in \mathcal{F}$.

For each n, let $G_n \in \mathcal{G}$ and $f_n \in \mathcal{B}$. Since $\mathcal{G}$ is a P-filter, there is some $G_\infty \in \mathcal{G}$ that is a pseudo-intersection for the G_n. As noted above, there is also some $f_\infty \in \mathcal{B}$ such that $f_\infty <^* f_n$ for every n. Then $X(G_\infty, f_\infty)$ is a pseudo-intersection for $\{X(G_n, f_n) : n \in \mathbb{N}\}$. Since the $X(G, f)$ form a basis for $\mathcal{F}$, it follows that $\mathcal{F}$ is a P-filter.

Finally we show that $\mathcal{F}^*$ is nowhere dense. Using Stone duality, we find that $\mathcal{F}^*$ is nowhere dense if and only if $\mathcal{F}$ is unbounded in $\mathcal{P}(\omega)/\text{fin}$. Suppose the latter is not the case, and let A be a bound for $\mathcal{F}$.

For each n, let

$$\alpha(n) = \max\{\min\{a - b_n, b_{n+1} - a\} : a \in A \cap [b_n, b_{n+1}]\}.$$

If the range of α is unbounded, then there is some function $f \in \mathcal{B}$ such that $\alpha \not<^* f$. In this case, $X(\mathbb{N}, f)$ misses infinitely many elements of A, contradicting the assumption that A is a pseudo-intersection for $\mathcal{F}$.

Therefore α is bounded: there is some k such that $\alpha(n) < k$ for all n. Fix N large enough so that $b_{n+1} - b_n > 2k$ whenever $n > N$. Fix $f \in \mathcal{B}$ so that $k \leq f(n) \leq \lfloor (b_n - b_{n-1})/2 \rfloor$ for all $n > N$. Define

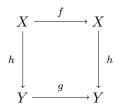$$B = \{n > N : |b_n - a| < k \text{ for some } a \in A\}.$$

Since $\mathcal{G}^*$ is nowhere dense, $\mathcal{G}$ has no pseudo-intersection; in particular, B is not a pseudo-intersection for $\mathcal{G}$. Thus there is some $G \in \mathcal{G}$ such that $B \setminus G$ is infinite. Then $X(G, f)$ misses infinitely many elements of A (namely all of those in $[b_n - k, b_n + k]$ for $n \in B \setminus G$), contradicting the assumption that A is a pseudo-intersection for $\mathcal{F}$. ∎

COROLLARY 4.9. *There is a nowhere dense closed $I \subseteq \mathbb{N}^*$ such that, for any $p, q \in \beta\mathbb{N}$, $p + q \in I$ if and only if $p \in I$.*

Proof. The proof is similar to that of Corollary 4.5(2). Let $I \subseteq \mathbb{N}^*$ be a nowhere dense P-set with $\sigma(I) = I$. Then I is a right ideal by Proposition 2.1, so if $p \in I$ then $p + q \in I$ for any q. If $q \notin I$, then $\{\sigma^n(p) : n \in \mathbb{N}\} \cap I = \emptyset$ because $\sigma(I) = I$; so $\overline{\{\sigma^n(p) : n \in \mathbb{N}\}} \cap I = \emptyset$ because I is a P-set. Thus $p + q = q\text{-}\lim_{n \in \mathbb{N}} \sigma^n(p) \notin I$. ∎

5. A consistent characterization of σ and σ^{-1}. In this section we show that it is both consistent with and independent of ZFC that the shift map σ and its inverse are, up to isomorphism, the unique chain transitive autohomeomorphisms of $\mathbb{N}^*$. The "independence" part of this result will use the nowhere dense P-set of Theorem 4.8.

Recall that two dynamical systems (X, f) and (Y, g) are *isomorphic* (or, for some authors, *conjugate*), written $(X, f) \cong (Y, g)$ or $X \cong Y$, if there is a homeomorphism $h : X \to Y$ such that the following diagram commutes:

$$
\begin{array}{ccc}
X & \xrightarrow{\ f\ } & X \\
\downarrow{\scriptstyle h} & & \downarrow{\scriptstyle h} \\
Y & \xrightarrow{\ g\ } & Y
\end{array}
$$

We say Y is a *quotient* of X (or is semiconjugate to X) if $h : X \to Y$ is a continuous surjection, but not necessarily a homeomorphism, for which

the above diagram commutes. In this case, h is called a *quotient map* (or a semiconjugation).

Chain transitivity is a well-studied property of metric dynamical systems. While it is typically defined in terms of some fixed metric, it is actually a topological property (i.e., it is preserved by isomorphisms), and, if one replaces metrics with uniformities, chain transitivity can be defined for non-metrizable systems as well.

Explicitly, if X is a dynamical system and $\mathcal{U}$ is an open cover of X, we say that $\langle x_i : i \leq n \rangle$ is a $\mathcal{U}$-*chain* from x_0 to x_n if, for every $i < n$, there is some $U \in \mathcal{U}$ such that $f(x_i) \in U$ and $x_{i+1} \in U$. Let us say that a (not necessarily metric) dynamical system X is *chain transitive* if for every open cover $\mathcal{U}$ of X and all $x, y \in X$, there is a $\mathcal{U}$-chain from x to y. The reader can easily check that this is equivalent to the usual definition of chain transitivity for metric systems (use the Lebesgue Covering Lemma for the nontrivial direction).

The following lemma is essentially proved in [A] as Theorem 4.12, although the notation is different and the proof there is for metric systems. For completeness, we give a proof here as well.

LEMMA 5.1. (X, f) *fails to be chain transitive if and only if there is some nonempty open* $U \neq X$ *and closed* $C \subseteq U$ *such that* $f(U) \subseteq C$.

Proof. Suppose there is some open $U \neq X$ and closed $C \subseteq U$ such that $f(U) \subseteq C$. Then $\mathcal{U} = \{U, X \setminus C\}$ is an open cover of X, and there is no $\mathcal{U}$-chain from a point of U to a point of $X \setminus U$.

For the other direction, suppose X is not chain transitive. Fix some open cover $\mathcal{U}$ of X, and some $x, y \in X$ with no $\mathcal{U}$-chain from x to y. Notice that if there is a $\mathcal{U}$-chain from x to any point of some $V \in \mathcal{U}$ then there is a $\mathcal{U}$-chain from x to every point of V. Let $\mathcal{U}_0$ be the set of all $V \in \mathcal{U}$ such that there is a $\mathcal{U}$-chain from x to some (any) point of V, and let $\mathcal{U}_1 = \mathcal{U} \setminus \mathcal{U}_0$. Set $U = \bigcup \mathcal{U}_0$ and $C = X \setminus \bigcup \mathcal{U}_1$. Then $\emptyset \neq U \neq X$ because $x \in U \not\ni y$. As $\mathcal{U}$ covers X, we have $C \subseteq U$. If $f(U) \setminus C \neq \emptyset$, then clearly it is possible to find a $\mathcal{U}$-chain from x to a point of $X \setminus C = \bigcup \mathcal{U}_1$, contradicting the definition of $\mathcal{U}_1$. Thus $f(U) \subseteq C$. ∎

LEMMA 5.2. *A map* $F : \mathbb{N}^* \to \mathbb{N}^*$ *is chain transitive if and only if, for every infinite, co-infinite* $A \subseteq \mathbb{N}$, $F(A^*) \not\subseteq A^*$.

Proof. Let U be an open subset of $\mathbb{N}^*$ with $U \neq \mathbb{N}^*$, and let $C \subseteq U$ be nonempty and closed. Since C is closed, $C = \bigcap \{A^* : A^* \supseteq C\}$. By [E, Corollary 3.1.5], there is some finite set $\{A_i : i \leq n\}$ such that $\bigcap_{i \leq n} A_i^* \subseteq U$. Since $\bigcap_{i \leq n} A_i^* = (\bigcap_{i \leq n} A_i)^*$, there is some $A \subseteq \mathbb{N}$ such that $C \subseteq A^* \subseteq U$. If $f(U) \subseteq C$ then $f(A^*) \subseteq A^*$, so the result follows from Lemma 5.1. ∎

THEOREM 5.3. *The shift map on $\mathbb{N}^*$ is chain transitive.*

Proof. Suppose A is infinite and co-infinite. Then $(A + 1) \cap (\mathbb{N} \setminus A) = \{n \in \mathbb{N} : n \in A \text{ and } n + 1 \notin A\}$ is infinite. In particular,

$$\sigma(A^*) \cap (\mathbb{N}^* \setminus A^*) = (A + 1)^* \cap (\mathbb{N} \setminus A)^* = ((A + 1) \cap (\mathbb{N} \setminus A))^* \neq \emptyset.$$

This shows that $\sigma(A^*) \nsubseteq A^*$. ∎

LEMMA 5.4. *If X is chain transitive and Y is a quotient of X, then Y is chain transitive.*

Again, this is proved in [A] for metric spaces, and the proof generalizes easily to the nonmetric case.

Proof of Lemma 5.4. Let $Q : X \to Y$ be a quotient map, let $\mathcal{U}$ be an open cover of Y, and let $y, z \in Y$. Then $\mathcal{V} = \{Q^{-1}(U) : U \in \mathcal{U}\}$ is an open cover of X. Fixing some $x_y \in Q^{-1}(y)$ and $x_z \in Q^{-1}(z)$, there is a $\mathcal{V}$-chain $\langle x_i : i \leq n \rangle$ in X from x_y to x_z. Then $\langle Q(x_i) : i \leq n \rangle$ is a $\mathcal{U}$-chain from y to z. ∎

We say that a map $F : \mathbb{N}^* \to \mathbb{N}^*$ is *trivial* if there is some map $f : \mathbb{N} \to \mathbb{N}$ such that $F = \beta f \restriction \mathbb{N}^*$.

LEMMA 5.5. *If τ is any chain transitive trivial autohomeomorphism of $\mathbb{N}^*$, then either $(\mathbb{N}^*, \tau)$ or $(\mathbb{N}^*, \tau^{-1})$ is isomorphic to $(\mathbb{N}^*, \sigma)$.*

Proof. Suppose $t : \mathbb{N} \to \mathbb{N}$ is a function such that $\tau = \beta t \restriction \mathbb{N}^*$ is a chain transitive autohomeomorphism of $\mathbb{N}^*$. By continuity, $\tau(A^*) = (t(A))^*$ for every $A \subseteq \mathbb{N}$. If $A \subseteq \mathbb{N}$ and $t(A) \subseteq^* A$, then $\tau(A^*) \subseteq A^*$. If $\emptyset \neq A^* \neq \mathbb{N}^*$, this contradicts the chain transitivity of τ by Lemma 5.2. Thus, unless A is either finite or cofinite, $t(A) \setminus A$ must be infinite.

Since βt is bijective, t restricts to an injection on a cofinite subset D of $\mathbb{N}$. By making t undefined on $\mathbb{N} \setminus D$, we may assume t is injective, hence invertible. For each n, let $\mathcal{O}(n) = \{t^m(n) : m \in \mathbb{Z}\}$. Note that the $\mathcal{O}(n)$ partition $\mathbb{N}$, that is, $n \in \mathcal{O}(m)$ if and only if $m \in \mathcal{O}(n)$.

Suppose there are infinitely many n with $\mathcal{O}(n)$ finite. Except for possibly the finitely many cases where $\mathcal{O}(n)$ contains an element of $\mathbb{N} \setminus D$ or $\mathbb{N} \setminus t(D)$, if $\mathcal{O}(n)$ is finite, then n is periodic under t and $\mathcal{O}(n)$ is invariant under t. If $\mathcal{A}$ is an infinite, co-infinite subset of $\{\mathcal{O}(n) : |\mathcal{O}(n)| < \aleph_0\}$, then $\bigcup \mathcal{A}$ is infinite and $t(A) \setminus A$ is finite. This contradicts the conclusion of the first paragraph of our proof, so there are finitely many n with $\mathcal{O}(n)$ finite.

In particular, $\mathcal{O}(n)$ is infinite for some n. Since $t(\mathcal{O}(n)) = \mathcal{O}(n)$, $\mathcal{O}(n)$ is cofinite, again using the first paragraph of our proof.

Fix n with $\mathcal{O}(n)$ cofinite. Clearly, either $\mathcal{O}^+(n) = \{t^m(n) : m \geq 0\}$ or $\mathcal{O}^-(n) = \{t^m(n) : m \leq 0\}$ is infinite. We suppose $\mathcal{O}^+(n)$ is infinite and show $(\mathbb{N}^*, \tau) \cong (\mathbb{N}^*, \sigma)$ (an analogous argument shows $(\mathbb{N}^*, \tau^{-1}) \cong (\mathbb{N}^*, \sigma)$ if $\mathcal{O}^-(n)$ is infinite).

Observe that $\mathcal{O}^+(n)$ is cofinite: otherwise $\emptyset \neq \mathcal{O}^+(n)^* \neq \mathbb{N}^*$ and $\tau(\mathcal{O}^+(n))^*$ $= \mathcal{O}^+(n)^*$. Let $\varphi : \mathbb{N} \to \mathbb{N}$ be given by $\varphi(m) = t^m(n)$. Consider $\beta\varphi :$ $\beta\mathbb{N} \to \beta\mathbb{N}$. Since φ is a bijection of $\mathbb{N}$ onto a cofinite subset of $\mathbb{N}$, $\beta\varphi$ restricts to an autohomeomorphism $\varphi^* : \mathbb{N}^* \to \mathbb{N}^*$. For every $A \subseteq \mathbb{N}$, we have $\varphi(A+1) = t(\varphi(A))$. By continuity, $\varphi^*(\sigma(A^*)) = \varphi^*((A+1)^*) = \tau(\varphi^*(A))$ for every A. It follows that, for every $p \in \mathbb{N}^*$, $\varphi^*(\sigma(p)) = \tau(\varphi^*(p))$. Thus φ^* is an isomorphism $(\mathbb{N}^*, \sigma) \to (\mathbb{N}^*, \tau)$. ∎

Notice that $(\mathbb{N}^*, \sigma) \cong (\mathbb{N}^*, \tau^{-1})$ if and only if $(\mathbb{N}^*, \sigma^{-1}) \cong (\mathbb{N}^*, \tau)$. We now have a consistent characterization of σ and σ^{-1}:

THEOREM 5.6. *It is consistent with ZFC that (up to isomorphism) σ and σ^{-1} are the only chain transitive autohomeomorphisms of $\mathbb{N}^*$.*

Proof. Shelah [S1] proves that it is consistent with ZFC that every autohomeomorphism of $\mathbb{N}^*$ is trivial. ∎

The next theorem gives the complementary result:

THEOREM 5.7. *Assuming the Continuum Hypothesis, there is a chain transitive autohomeomorphism of $\mathbb{N}^*$ not isomorphic to σ or σ^{-1}.*

Proof. Using Theorem 4.8 (or Theorem 3.5 plus CH), let $X \subseteq \mathbb{N}^*$ be a nowhere dense P-set with $\sigma(X) = X$. Let Y be the quotient space obtained from $\mathbb{N}^*$ from the relation

$$p \sim q \iff p = q \text{ or } p, q \in X.$$

That is, Y is the compact Hausdorff space obtained from $\mathbb{N}^*$ by collapsing X to a point.

Because $\sigma(X) = X$, $[\sigma(x)] = [\sigma(y)]$ whenever $x \sim y$; define τ on Y by setting $\tau([x]) = [\sigma(x)]$. It is easy to check that this map is continuous, so that (Y, τ) is a dynamical system. Clearly, τ is a bijection on Y (because σ is a bijection on $\mathbb{N}^*$), so that τ is a continuous bijection, hence an autohomeomorphism, on Y. By construction (Y, τ) is a quotient of $(\mathbb{N}^*, \sigma)$ and is therefore chain transitive by Lemma 5.4.

Because X is nowhere dense in $\mathbb{N}^*$, Corollary 1.2.4 in [vM] implies that $Y \approx \mathbb{N}^*$. Thus τ is a chain transitive autohomeomorphism of $\mathbb{N}^*$. We note that Corollary 1.2.4 in [vM] requires CH; in fact, it has been shown that a version of it is equivalent to CH (see [vDvM]).

To complete the proof, we observe that $\tau([x]) = [x]$ for $x \in X$. Since σ has no fixed points, (Y, τ) can be isomorphic to neither $(\mathbb{N}^*, \sigma)$ nor $(\mathbb{N}^*, \sigma^{-1})$. ∎

Roughly speaking, there are many models in which CH fails and not every autohomeomorphism of $\mathbb{N}^*$ is trivial. We leave open the question of whether some of these contain chain transitive autohomeomorphisms of $\mathbb{N}^*$ other than σ and σ^{-1}. We also leave open the question of whether another

adjective or two can be added to "chain transitive autohomeomorphism" to obtain a topological characterization of σ in ZFC.

References

[A]   E. Akin, *The General Topology of Dynamical Systems*, Grad. Stud. Math. 1, Amer. Math. Soc., Providence, RI, reprint ed., 2010.

[Be]   V. Bergelson, *Minimal idempotents and ergodic Ramsey theory*, in: Topics in Dynamics and Ergodic Theory, London Math. Soc. Lecture Note Ser. 310, Cambridge Univ. Press, Cambridge, 2003, 8–39.

[BH]   V. Bergelson and N. Hindman, *Nonmetrizable topological dynamics and Ramsey theory*, Trans. Amer. Math. Soc. 320 (1990), 293–320.

[BHM] V. Bergelson, N. Hindman, and R. McCutcheon, *Notions of size and combinatorial properties of quotient sets in semigroups*, Topology Proc. 23 (1998), 23–60.

[Bl1] A. Blass, *Ultrafilters: where topological dynamics = algebra = combinatorics*, Topology Proc. 18 (1993), 33–56.

[Bl2] A. Blass, *Combinatorial cardinal characteristics of the continuum*, in: Handbook of Set Theory, A. Kanamori et al. (eds.), Springer, Dordrecht, 2010.

[BlH] A. Blass and N. Hindman, *On strongly summable ultrafilters and union ultrafilters*, Trans. Amer. Math. Soc. 304 (1987), 83–99.

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

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

[vDvM] E. K. van Douwen and J. van Mill, *Parovičenko's characterization of $\beta\omega - \omega$ implies CH*, Proc. Amer. Math. Soc. 72 (1978), 539–541.

[E]   R. Engelking, *General Topology*, 2nd ed., Sigma Ser. Pure Math. 6, Heldermann, Berlin, 1989.

[HS]   N. Hindman and D. Strauss, *Algebra in the Stone–Čech Compactification*, de Gruyter, Berlin, 1998.

[K]   K. Kunen, *Weak P-point in $\boldsymbol{N}^*$*, in: Topology, Colloq. Math. Soc. János Bolyai 23, North-Holland, Amsterdam, 1980, 741–749.

[MS]   M. Malliaris and S. Shelah, *Cofinality spectrum theorems in model theory, set theory and general topology*, arXiv:1208.5424 (2012).

[vM]   J. van Mill, *An introduction to $\beta\omega$*, in: Handbook of Set-Theoretic Topology, K. Kunen and J. E. Vaughan (eds.), North-Holland, Amsterdam, 1984, 503–567.

[R]   W. Rudin, *Homogeneity problems in the theory of Čech compactifications*, Duke Math. J. 23 (1956), 409–419, 633.

[S1]   S. Shelah, *Proper Forcing*, Lecture Notes in Math. 940, Springer, Berlin, 1982.

[S2]   S. Shelah, *Proper and Improper Forcing*, Perspectives Math. Logic, Springer, Berlin, 1998.

[SS]   S. Shelah and J. Steprāns, *PFA implies all automorphisms are trivial*, Proc. Amer. Math. Soc. 104 (1988), 1220–1225.

[W]   E. L. Wimmers, *The Shelah P-point independence theorem*, Israel J. Math. 43 (1982), 28–48.

[Z]   Y. Zelenyuk, *Principal left ideals of βG may be both minimal and maximal*, Bull. London Math. Soc. 45 (2013), 613–617.

W. R. Brian
Department of Mathematics
Tulane University
6823 St. Charles Ave.
New Orleans, LA 70118, U.S.A.
E-mail: wbrian.math@gmail.com