Pittsburgh Interdisciplinary Mathematics Review **5** (2026), 13–40

When a set theorist hears "combinatorics", Part 2: Trees

Thomas Gilton

(Communicated by Joe Tracy
Copyedited by Lucas Pasdo)

Abstract

In the previous installment of this series (When a Set Theorist Hears "Combinatorics": Ramsey Theory, [Gil25]) we discussed the combinatorial problem of finding large patches of order in graphs, and this was done from the perspective of a set theorist. Thus we focused on infinite graphs, both countably infinite and uncountably infinite. In this follow-up expository article, we focus on trees, another common object in combinatorics. Working yet again from the perspective of a set theorist, we will focus on infinite trees. We shall begin with graph-theoretic trees to set the stage. After subsequently developing the tools to talk about trees with a "longer than infinite" height, we will consider several kinds of uncountable trees, how they structurally differ from countably infinite trees, and how they are related to a question about characterizing the real numbers.

1. Introduction

In the previous paper ([Gil25]) we took a look at Ramsey theory, but from the perspective of a set theorist. Generally speaking, Ramsey theory concerns itself with finding large patches of order (cliques or independent sets) in otherwise quite messy objects (graphs). Ramsey theory is an active area of research in both finite and infinite combinatorics, but to a set theorist, the case of infinite graphs is the focus.

In this paper, we do something similar, but instead of Ramsey theory, we shall focus on trees. Just as in the previous paper, the object of study is deeply important and fascinating in both the finite and infinite cases. However, we will focus on the case of infinite trees, since that is where the set-theoretic techniques really shine.

We will kick off the paper in Section 2 by briefly reviewing trees in the graph-theoretic setting. The first key take-away will be how to extract a partial order from a tree. Then we will show that if such a (graph-theoretic) tree is very tall, but not too wide, then it has an infinite path going up from the root (a "branch", in the jargon). At this level, we don't really need the set-theoretic techniques.

However, beginning in Section 3, we will develop the ideas necessary to talk about trees that grow "higher than infinitely high" (more precisely, trees whose height is an uncountable ordinal). We will develop the basic ideas of well-orders and their canonical representations, called ordinals (lengths of higher infinite sequences).

Having done this spade work, we transition in Section 4 to discussing trees from the perspective of a set theorist. A key point is that we will generalize the order structure that we saw in Section 2 by using well-orders.

Sections 5 and 6 are devoted to applying these ideas. In the former section, we show that an uncountable tree that isn't too wide *can* fail to have a "cofinal branch"; thus the generalization of the main result in Section 2 fails once we toggle the parameters. Section 6 is devoted to connecting a question about what uniquely characterizes the reals $\mathbb{R}$ to the question of the existence of a certain kind of tree.

The reader who is not familiar with set-theoretic techniques is invited to read the paper at least twice: first just noting the main ideas, and only diving into the proofs on a second read.

2. Graph-theoretic trees and König's theorem

Let's begin by recalling what a tree is in the context of graph theory:

Definition 2.1. A tree is a graph $G = (V, E)$ so that

1. G is connected and
2. G has no cycles.

It would be difficult to overstate just how important graph-theoretic trees are. Pick up a book on combinatorics and discrete math if you're not convinced.

To set the stage for later work, let's describe **how to generate a partial order from a tree** (see the picture below the next two paragraphs).

From a given graph-theoretic tree, you can isolate a given node as the "root" and then proceed to define the levels of the tree. A crucial fact is that, because we have a tree, there is a *unique path* from a given node to the root. Thus we may sensibly say that level 1 consists of all nodes that have an edge to the root; level 2 consists of all nodes that are a distance of exactly 2 away from the root; and level n consists of

all nodes that are a distance of exactly n from the root. And given that there is a unique path from a node down to the root, no node is on multiple levels.

You can then define an ordering on this tree by saying that node u is less than node v if u is connected to v but on a lower level than v. Let's make two observations that will be crucial for our subsequent discussion of higher infinite trees:

1. This ordering on the tree is a *partial order*: if you take two nodes at the same level, for example, they will not be comparable.

2. Additionally, if v is a node in the tree, then the set of nodes below it, i.e., $\{u \in V : u < v\}$ is a *finite* linear order! Though it might not make sense yet, make sure to read the following sentence: **every finite linear order is a well-order**. That will be important later.

The following picture illustrates these ideas:

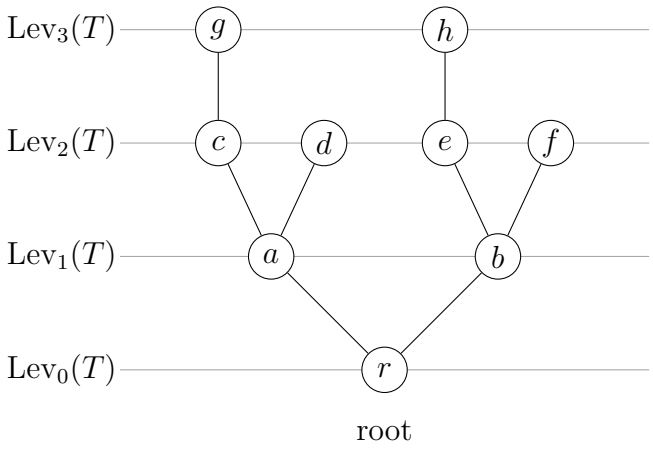

Note, for example, that $r < c < g$ and $r < d$, but neither $c < d$ nor $d < c$. Also note, for example, that level 2 has four nodes and level 3 has two nodes. Finally, notice that g together with the set of nodes below g, namely $\{r, a, c\}$, forms a linear order.

You can think of (finite) trees, therefore, as a bunch of finite linear orders joined together, with lots of splitting. The linear orders are the same as paths (without repeats) coming up from the root. We call these paths *branches*.

Let's move on to infinite trees! And here we're interested in which infinite trees have these infinite linearly-ordered subsets. Or in other words, which infinite trees have infinite branches. Let's make this more precise in the following question:

Question 2.2. Say we have an infinite tree T, in the graph-theoretic sense. Can we find an infinite, linearly-ordered subset of T? This is the same as finding a sequence of nodes $v_0 < v_1 < v_2 < ...$ in the tree, starting from the root, where v_k is on level k.

Such a path would be an *infinite branch*.

This question has a lovely answer: if the tree is infinite, but *each of the levels are finite*, then yes, we can find an infinite branch (i.e., an infinite linearly-ordered subset). Before we give the proof of this, let's first look at an example which illustrates what can go wrong if the levels are infinite. Consider the tree from the following picture in which, for each $n \geq 1$, there is a distinct branch of length n coming up from the root. Notice that while the tree has, for each $n \geq 1$, a branch of length n, it doesn't have any infinite branch (i.e., an infinite, linearly- ordered subset). And here's the problem: each level is infinite!

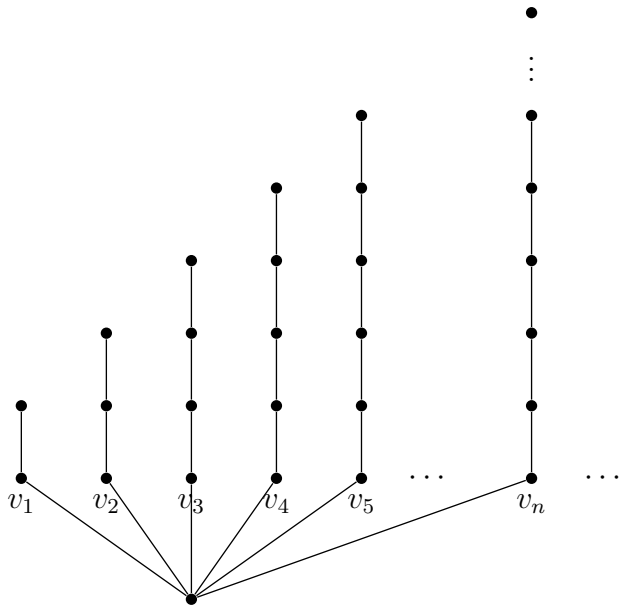

Now that we see what can go wrong when the tree is too wide, let's see how having finite levels guarantees an infinite branch. This is due to König ([Kön27]) from a 1927 paper:

Theorem 2.3. *(König, 1927) If T is an infinite tree all of whose levels are finite, then T has an infinite path.*

Proof. We will construct the branch $v_0 < v_1 < v_2 < \dots$ by recursively picking a node level-by-level. To ensure we don't run out of nodes (i.e., to ensure that we don't run into a "leaf"), we maintain as a recursion hypothesis that every node that we pick has infinitely-many nodes above it in the tree. *This will be an exercise in the pigeonhole principle.*

Let v_0 be the root of T. (That was easy.) We now describe the construction of v_1 somewhat loosely so the reader can get a feel for things, and then we do the general case more precisely.

By assumption, there are only finitely-many nodes on level 1. Let these nodes be enumerated as $v_{1,1}, \ldots, v_{1,k_1}$. We apply the pigeonhole principle: every node in T at a higher level is above one of the nodes on level 1. But there are infinitely-many such nodes. Thus there must be at least one node on level 1, say $v_{1,j}$, which has infinitely-many nodes above it. Set v_1 to equal $v_{1,j}$.

Now suppose that we have defined $v_0, v_1, \ldots, v_n$ so that

1. $v_0 < v_1 < \cdots < v_n$ form a path with v_k on level k, and

2. there are infinitely-many nodes of T which are above v_n. That is to say, the set $\text{Above}(v_n) := \{v \in T : v_n < v\}$ is infinite.

To define v_{n+1}, we will apply the pigeonhole principle to see that, among the nodes of T on level $n+1$ which are above v_n, at least one of them must have infinitely-many nodes of T extending it. To make this precise, note that every $v \in \text{Above}(v_n)$ is either on level $n+1$ or on some higher level. Let $\text{Next}(v_n)$ be all the nodes in $\text{Above}(v_n)$ which are on level $n+1$. Since each level of T is finite, and since $\text{Next}(v_n)$ consists of nodes entirely on level $n+1$, it is finite too. Now define $B_n := \text{Above}(v_n) \setminus \text{Next}(v_n)$, so that B_n consists of all nodes above v_n which are on a higher level than $n+1$. Note that B_n is infinite.

Now let's set up a pigeonhole argument: each $v \in B_n$ has a *unique* node below it on level $n+1$, because T is a tree. Thus we can partition the $v \in B_n$ into finitely many pieces based upon which node of $\text{Next}(v_n)$ sits below it. This is a partition of an infinite set into finitely-many pieces, and therefore at least one piece is infinite. Consequently, there is at least one node $v_{n+1} \in \text{Next}(v_n)$ with infinitely-many nodes above it.

This completes the construction of the infinite branch $v_0 < v_1 < v_2 < \ldots$ through T and thereby the proof of the theorem. □

The pigeonhole principle, again.

Remark 2.4. *Let's summarize this section: from a given graph-theoretic tree, and a given choice of a root, you can define (i) levels and (ii) a partial order. While the whole order is only a partial order, the set of predecessors of a given node is a linear order. In fact, it is a* ***finite linear order****, and all finite linear orders are* ***well-orders*** *(this will make more sense later).*

Given that linearly-ordered subsets (i.e., branches) of a tree are "nice", we want to know when infinite linearly-ordered subsets of an infinite tree exist. A key take-away is that for the trees in Theorem 2.3, "not being too wide" implies that "an infinite branch exists."

We will see in a later section that the relevant generalization is ***not*** *true once we hit uncountable trees (at least for so-called ω_1-trees).*

3. Basic set-theoretic ideas needed for taller trees

In this section, we will develop the vocabulary and basic ideas to talk about uncountable trees. To get an idea of why it would be helpful to do so, try to think about how you'd define a tree that has (loosely speaking) and "infinity-eth" level, i.e., a level that is "higher" than all the finite levels.

The trees from the previous section all had levels at a finite height, since each node was connected (in the graph-theoretic sense) to the root. Or put differently, the path from a node to the root, *being a path connecting two nodes*, is finite. What would it even mean to have a node that is infinitely high up on the tree? In short,

what would it mean to have a "longer than infinite" path?

In this section, we're going to make these ideas precise.

To set the stage, a tree (to a set theorist) will be a type of partial order $(T, <)$ (recall that we discussed in the previous section how to turn a graph-theoretic tree into a partial order). We will require that for every $y \in T$, the set of nodes in T below y, namely

$$\{x \in T : x < y\},$$

forms a very nice kind of linearly-ordered set called a *well-order*. **Notice how this is analogous to the previous section:** there we extracted a partial order from a tree, and the set of predecessors of a node v (i.e., the nodes on the unique path from v to the root) formed a *finite* linear order. And as remarked earlier, every finite linear order is a well-order. This is how we're going to generalize the ideas from the previous section to the context of higher infinities.

To ensure that we are on the same page, we will need to discuss what well-ordered sets are. This will lead us to discuss the canonical well-orders known as the ordinals. Here "canonical" means that we take these ordinals as the standard and compare every other well-order to them. Among ordinals, we will find cardinals, which are size-numbers.

Armed with these notions, we can then define more precisely what a tree is, as well as other related and important vocabulary.

3.1. **Well orders and limit elements**

To begin this subsection, here's the definition of a well-order:

Definition 3.1. A linear order $<$ on a set L is called a *well-order* if every non-empty $Z \subseteq L$ has a $<$-least element.

You are already familiar with these: this is one of the central features of the natural numbers, $\mathbb{N}$! If you have a set S of natural numbers, and S is non-empty, then S has a least element. This is precisely what makes induction on $\mathbb{N}$ a sensible proof strategy, and it's what makes it rational to define objects by recursion.

But there's nothing that says you can *only* do recursion and induction on $\mathbb{N}$. You can do them on *any* well-order, and for pretty much the same reason as on $\mathbb{N}$. Let's begin by proving a basic fact about well-orders: if an element is not the maximum in the order, then there is a least element bigger than it:

Lemma 3.2. *Suppose that $(L, <)$ is a well-order, that $x \in L$, and that x is not the*

maximum element. Then x has an immediate successor, i.e., there is a $<$-least $y \in L$ which is above x.

Proof. By assumption, $\{z \in L : x < z\}$ is non-empty. Apply the definition of a well-order and let y be the minimum element. □

This allows us to make the following definition:

Definition 3.3. Suppose that $(L, <)$ is a well-order, that $x \in L$, and that x is not the maximum element. We use "$x+1$" to abbreviate the successor of x in L, namely, the least $y \in L$ with $x < y$.

Now you might initially think that if $(L, <)$ is a well-order, then we can do induction on it just like on $\mathbb{N}$. And this is almost true...The issue is that, because $\mathbb{N}$ is not a very "long" well-order, there is an additional complexity that occurs in more general well-orders. Recall that to do induction on $\mathbb{N}$, you need to prove that the statement of interest, say P, holds at 0 and that for all n, *if* $P(n)$ holds, *then* $P(n+1)$ holds. And this makes good sense: notice crucially that every $n \in \mathbb{N}$ is either 0 or is the successor of something (for example, 3 is the successor of 2). There are no other cases.

By contrast, set-theorists frequently use much longer well-orders. And these longer well-orders have an additional, third case (not just base case and successor). To see this, let's think about how we could make a "longer" well-order than that of the natural numbers. We'd need to find some element, call it ∞, and declare that ∞ is bigger than each natural number. So here our underlying set is $\mathbb{N} \cup \{\infty\}$, and the ordering $\preceq$ is the usual ordering when restricted to $\mathbb{N}$, and otherwise it says $n \prec \infty$. Here's a picture:

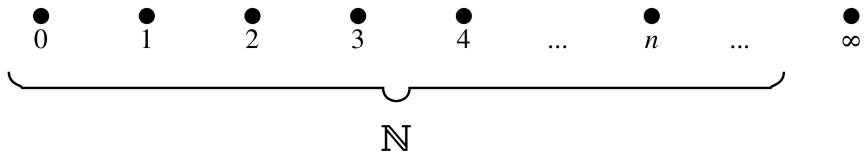

Now this set with a "top" element is still a well-order: let $A \subseteq \mathbb{N} \cup \{\infty\}$ be non-empty. We have some cases: if A contains a natural number, then it contains a least such, and this is in fact the least element of A. Otherwise, A would have to be just the singleton $\{\infty\}$.

Now what's different here? *The element ∞ does not have a predecessor,* or put differently, ∞ is not the successor of anything. If $n \in \mathbb{N}$, then n is not the predecessor

of ∞ since the next element, $n+1$, is still below ∞. We refer to such elements in a well-order as limit elements:

Definition 3.4. Let $(L, <)$ be a well-order. An element $x \in L$ is said to be a *limit element* if x is not the least element of L and if x has no immediate predecessors (that is, if whenever $y < x$ then $y + 1 < x$ too).

We now see that these exhaust all of the possible cases: given a well-order $(L, <)$ and $x \in L$, if x is not the least element of L, then either (a) x has an immediate predecessor or (b) x does not (in which case it's a limit element). Thus we can still make induction and recursion run on an arbitrary well-order, but in general, we must deal with this third case. Here's the precise statement:

Theorem 3.5. *Suppose that $(L, <)$ is a well-order, that $P(u)$ is a statement, and that the following three conditions hold:*

(base case) *$P(0_L)$ is true (where 0_L is the least element of L).*

(successor case) *If $x \in L$ and $x + 1$ exists (i.e., x is not the top element), and if $P(x)$ is true, then $P(x + 1)$ is true.*

(limit case) *If x is a limit element and if for all $y < x$ $P(y)$ is true, then $P(x)$ is true.*

Then $P(x)$ is true for all $x \in L$.

Proof. Suppose otherwise. Then $\{x \in L : P(x) \text{ is false}\}$ is non-empty and hence has a least element, call it z. Now z cannot be 0_L, since $P(0_L)$ is true. Thus z is either a limit element or a successor element. Since z is the least where P is false, we know that $P(x)$ holds for all $x < z$. We now get a contradiction by taking cases on whether z is equal to $\bar{z} + 1$ for some $\bar{z}$ or whether z is a limit element. □

The upshot: you can prove statements by doing induction along a well-ordered set, and you can define objects by recursion on a well-ordered set.

3.2. **Canonical well-orders: ordinals**

One of the most lovely features of well-orders is that there are simple, canonical well-orders. These are known as the ordinals. And for a set theorist, these ordinals are the "lengths" of a given well-order. We're going to give the definition of an ordinal, some examples, and then we'll summarize the key things you need to know about them for the rest of this paper. We won't be working directly with the definition to prove these facts about ordinals, so don't worry if the following definition seems opaque; just keep reading.

Definition 3.6. A set α is an *ordinal* if the set-theoretic membership relation $\in$ is a well-order on α, and if for every $\beta \in \alpha$, β is also a subset of α.[1]

Notice how simple this is: when doing set theory, the only thing you're guaranteed is a $\in$ relation, and we're using that simplest of tools to define orders. Now here are some crucial facts about ordinals:

Proposition 3.7. *Let α and β be ordinals.*

1. *If $(\alpha, \in)$ is isomorphic as a well-order to $(\beta, \in)$, then in fact $\alpha = \beta$ (they are literally the same set).*
2. *If $\alpha \neq \beta$, then either $\alpha \in \beta$ or $\beta \in \alpha$.*
3. *If $(L, <)$ is a well-order, then there is a unique ordinal δ so that $(L, <)$ is isomorphic to $(\delta, \in)$.*
4. *$\alpha \cup \{\alpha\}$ is the least ordinal greater than α. We denote this ordinal by $\alpha + 1$.*

The upshot of this is that it's substantially easier to work with ordinals than with arbitrary well-orders. This is because the actual ordering on an ordinal is very simple: it's just membership! Going forwards, you want to think of ordinal numbers as giving (usually infinite) lengths. Now for some examples.

The smallest ordinal is the empty set, $\emptyset$, and when doing set theory, we prefer to use the notation 0. The next ordinal is 1, which we think of as $\{0\}$. Then comes "2", which as an ordinal is equal to $\{0, 1\}$. And so on for all the natural numbers. So for example, $42 = \{0, 1, 2, \ldots, 40, 41\}$.

Now we can generate the next ordinal above all of these by collecting them into one new set. We define ω to be the ordinal whose elements are all the natural numbers...in other words, the ordinal ω is just equal to $\mathbb{N}$! (This is weird the first time you read it: you're probably not used to thinking of the set of natural numbers as itself a number.) ω is called a *limit ordinal* since it does not have an immediate predecessor.

At this point, we can define the next ordinal, $\omega + 1$, and this equals $\omega \cup \{\omega\}$ (see part (4) of Proposition 3.7). Then comes $\omega + 2 = (\omega + 1) \cup \{\omega + 1\}$ and so on. The next limit ordinal after ω is $\omega + \omega$, typically denoted[2] $\omega \cdot 2$. Here's a picture:

[1] In other words, every element of an element of α is itself an element of α. In the jargon, we abbreviate this by saying that $\in$ is "transitive" on α. We won't need this idea explicitly in the paper, but it is crucial for establishing the basic properties of the ordinals that we use.

[2] This is not the same thing as $2 \cdot \omega$. The way ordinal multiplication is defined, $2 \cdot \omega$ just equals ω itself. This is because you're thinking of it as infinitely-many copies of the ordinal 2 stacked on top of each other.

So for $\omega + \omega$ you have your first copy of ω (i.e., the natural numbers), and then you start another copy "after" the first one.

3.3. **A brief word on cardinal numbers**

In this subsection, we will briefly define cardinal numbers and then use all of this in the next sections to start talking about trees.

We've seen that ordinals represent well-ordered lengths, especially infinite ones. Cardinals (which are ordinals with additional properties) represent sizes, finite and infinite. Here's the definition of a cardinal number:

Definition 3.8. An ordinal κ is said to be a *cardinal* if κ is not in bijection with any $\alpha \in \kappa$ (i.e., not in bijection with any smaller ordinal).

So an ordinal number κ is a cardinal if it has a greater cardinality than all of its predecessors. This definition coincides with your usual understanding of size for finite objects: the natural numbers 0, 1, 2, 3, etc. are all cardinals as well as ordinals. Moreover, the first infinite ordinal, ω, is also a cardinal, since it's not in bijection with any finite n; note that countably infinite sets are exactly those in bijection with ω.

But at precisely this point, we see that ordinals (lengths) and cardinals (sizes) start to diverge. The next ordinal after ω, $\omega + 1$, is not a cardinal since $\omega + 1$ is in bijection with ω. In fact the function

$$f(x) = \begin{cases} 0 & \text{if } x = \omega \\ n+1 & \text{if } x = n \end{cases}$$

is a bijection from $\omega + 1 = \omega \cup \{\omega\}$ to ω (you just scoot the top element ω to the bottom and bump everyone else up by 1).

As another example, $\omega + \omega$ is not a cardinal. Roughly, we'll take the first block and biject it with the even natural numbers and the second block and biject it with the odds. More precisely (and recall that elements of $\omega + \omega$ are either naturals n or of the form $\omega + n$ for a natural n) the following function will do the trick:

$$f(x) = \begin{cases} 2x & \text{if } x \in \omega \\ 2n+1 & \text{if } x = \omega + n \end{cases}$$

We now have a rough sense of what cardinals are. Let's now introduce some notation for the next two cardinals after ω (that's all we'll need for this paper):

Definition 3.9. ω_1 is the least cardinal after ω, i.e., the least uncountable cardinal. ω_2 is the least cardinal after ω_1.

At this point we finally have enough set-theoretic terminology to talk about trees. That's the focus of the next section.

4. **Set-theoretic trees**

As discussed in the previous section, we want to generalize the notion of how high up a node can be in a tree. We'll get a more general definition that reduces to the definitions in the graph-theoretic context in all of the cases that you're used to. We're going to keep the idea of levels of a tree, and the idea of an ordering, but generalize the idea of a path and the length of a path by focusing almost entirely on the order structure.

Here well-orders *and hence ordinals*, are the key to moving forwards. A tree for us is going to be a kind of partial order. We're going to require that if y is a node in the (set-theoretic) tree, then the set of predecessors of y is well-ordered. So there's still, in a sense, a unique way to get down to the root from y, and this "path" (if you will)[3] is thought of as a well-order, perhaps a very long one!

So here's the definition, finally:

Definition 4.1. A *tree* is a partially ordered set $(T, \leq)$ so that for any $y \in T$, the set $\{x \in T : x < y\}$ is well-ordered by $\leq$.

So you can think of a tree as a bunch of ordinals branching off of each other.

Associated to a tree $(T, \leq)$, we have notions similar to those in the graph-theoretic case:

Definition 4.2. Let $(T, \leq)$ be a tree.

1. The *height of a node* y is the length of the well-order $\{x \in T : x < y\}$, i.e., the unique ordinal α which is isomorphic to that well-order.
2. The *height of* T is $\sup\{\text{height}(y) + 1 : y \in T\}$.[4]

[3]Notice the scare quotes: these things are not actually paths. They generalize the way in which we use paths in graph-theoretic trees.

[4]Don't worry too much about the extra "+1" in the definition. It's there to help us avoid some silliness.

3. For each $\alpha < \text{height}(T)$, *the αth level of T* consists of all $y \in T$ so that y has height α. We denote this with $\text{Lev}_\alpha(T)$.

4. A *branch* is any well-ordered subset of T, and a branch b is said to be *cofinal* if b contains a node on each level of T.

Let's look at some examples!

Example 4.3. Let T consist of all finite binary sequences. Here are some examples of members of T, the empty sequence $\emptyset$; the two sequences of length one, $\langle 0\rangle$ and $\langle 1\rangle$; and the sequence $\langle 1, 0, 0, 1, 1\rangle$. We use $2^{<\omega}$ to denote the set of all finite binary sequences, viewing each such sequence as a function from a natural (i.e., some $n < \omega$) to $2 = \{0, 1\}$ (recall our discussion of ordinals!).

We define the ordering by end-extension: if s and t are finite binary sequences, $s \sqsubset t$ if t restricted to the length of s equals s. For example, $\langle 0, 1, 1\rangle \sqsubset \langle 0, 1, 1, 0, 0\rangle$ but $\langle 0, 1, 1\rangle \not\sqsubset \langle 1, 1, 1, 1\rangle$.

The nth level of this tree consists of all binary sequences of length n (so that the height of a sequence in the tree is the length of the sequence). Since there are nodes of height n for each n, but no nodes of height ω (i.e., no infinite sequences are members of the tree), the height of T is ω.

Any function $f : \omega \to 2 = \{0, 1\}$ is an infinite branch through the tree! In fact, the set of infinite branches through this tree is homeomorphic to the Cantor set...

Example 4.4. We now modify the previous example by allowing for sequences of much greater length. So, let T be the set of all functions $f : \alpha \to \{0, 1\}$, where $\alpha < \omega_1$. That is, the elements of T are exactly the binary sequences whose length is a countable ordinal. As before, the height of a node is the length of the sequence, and consequently, the height of T is ω_1. The α-th level of T consists of all $f : \alpha \to \{0, 1\}$ (all binary sequences of length α). And finally, the cofinal branches through T are precisely the functions $f : \omega_1 \to 2$.

There is one very interesting difference between this example and the previous one, and it concerns how wide the trees are. In the previous example, each level was finite (more precisely, level n had size 2^n). In particular, each level of the tree had a size that is much smaller than the height of the tree.

By contrast, in the second example, once α is an infinite ordinal, level α has size continuum (recall that this equals the size of the powerset of the naturals). For example, consider level ω, and note that *every* function $f : \omega \to 2$ is on level ω. Note also that there are continuum-many such functions. Now the continuum is *at least* size ω_1 (the CH says that it's exactly ω_1). Thus for this second example, the tree is at least as wide as it is tall.

Going forwards, we will want to exclude from our investigation trees which are too wide, such as the tree from the second example. And this brings us to out next section on ω_1-trees.

5. Trees on ω_1 and the existence of ω_1-Aronszajn trees

In this section, we will consider trees that have height ω_1 but levels that are at most countable. Recall from Theorem 2.3 that if T has height ω but each level is finite, then T has a cofinal branch. The motivating question for us is whether this generalizes when we "bump up" both of those parameters to higher cardinals:

Question 5.1. Does every tree T of height ω_1 all of whose levels are at most countable have a cofinal branch (namely, a well-ordered subset of length ω_1)?

Put differently, do tall, skinny trees on ω_1 still have cofinal branches? Since we talk about such trees so often, it is helpful to give them a name:

Definition 5.2. An *ω_1-tree* is a tree T so that T has height ω_1 and for each $\alpha < \omega_1$, $\mathrm{Lev}_\alpha(T)$ is at most countable.

A tree that would be a counterexample to Question 5.1 is given a special name (pronounced Erin-shine):

Definition 5.3. An *ω_1-Aronszajn* ("Erin-shine") *tree* is an ω_1-tree without a cofinal branch.

So Question 5.1 asks: do any ω_1-Aronszajn trees exist? And it turns out that the answer is yes! We can prove, with the usual ZFC axioms of set theory, that an ω_1-Aronszajn tree exists. The construction is due, naturally enough, to Aronszajn, but it appeared in Kurepa's 1935 paper [Kur35]). We'll give a proof after one more definition (which we'll need in the proof).

Definition 5.4. Let X be an infinite set and f, g two functions with domain X. We say that f and g are *equal almost everywhere*, denoted $f =^* g$, if $\{x \in X : f(x) \neq g(x)\}$ is finite.

Now for the theorem (note that the proof is fairly involved; I recommend reading it only on a second pass through the paper):

Theorem 5.5. *There exists an ω_1-Aronszajn tree.*

Proof. We will construct an ω_1-Aronszajn tree T recursively, level-by-level. Here are the basic structural features: for each countable ordinal $\alpha < \omega_1$, $\mathrm{Lev}_\alpha(T)$ will consist of functions f from α to $\mathbb{N}$ (for now, any countably infinite set will do; later constructions use countable sets with additional features, such as $\mathbb{Q}$). Note that we can't include *all* such functions on level α, since there are uncountably-many such functions, and we need our levels to be countable. The ordering on the tree is as follows: given $f \in \mathrm{Lev}_\alpha(T)$ and $g \in \mathrm{Lev}_\beta(T)$ with $\alpha < \beta$, we say $f \sqsubset g$ exactly when $g \restriction \alpha = f$, i.e., when g end-extends f.

In addition to the above, we require that every function $f \in T$ is *injective*. Notice, crucially, that *IF* we can maintain this requirement throughout the construction of the tree T, then T will not have a cofinal branch. Indeed, any branch b through T gives rise to an injective function (just take the union of the functions along the branch) from an ordinal α into $\mathbb{N}$; and any set that injects into a countable set is itself countable.

Recall that we need to also make sure that the tree has height ω_1 (so that it is tall enough) and has levels that are at most countably infinite (so that the tree is not too wide). The following recursive assumptions will help us ensure that we don't run into any obstacles as we build the tree:

1. (Extendable) suppose that we've constructed T up to level γ, and let $\alpha < \beta < \gamma$. Suppose that $f \in \mathrm{Lev}_\alpha(T)$. Then there exists $g \in \mathrm{Lev}_\beta(T)$ with $f \sqsubset g$. This means that every node can be extended to every higher level in the tree.

2. (Closure) suppose that we've constructed T up to level γ, and let $\alpha < \gamma$. Let $f \in \mathrm{Lev}_\alpha(T)$. Then (recalling Definition 5.4) for any injective $h : \alpha \to \mathbb{N}$ with $f =^* h$, $h \in \mathrm{Lev}_\alpha(T)$ too. In other words, every injective function which is a finite modification of a function in T is also a function in T.

3. (Options) if f is a function in T, then $\mathbb{N} \setminus \mathrm{ran}(f)$ is infinite. In other words, not only is f injective, but there are infinitely-many values that f does *not* take.

Let's get started with the construction. For each finite n, the nth level of T consists of all injective $f : n \to \mathbb{N}$. Note that all of the recursive assumptions hold true here. This is the boring part.

The proof now divides into the successor and limit cases of the construction.

For the successor case, suppose that we've constructed $\mathrm{Lev}_\alpha(T)$ (and hence all previous levels...). To construct $\mathrm{Lev}_{\alpha+1}(T)$, let $f \in \mathrm{Lev}_\alpha(T)$ be arbitrary. For each $k \in \mathbb{N} \setminus \mathrm{ran}(f)$ (by the "options" recursion hypothesis, this set is infinite), we add k on top of f. More precisely, we add the function g to $\mathrm{Lev}_{\alpha+1}(T)$, where $g \restriction \alpha = f$ and where $g(\alpha) = k$. Note that g is injective and that $\mathbb{N} \setminus \mathrm{ran}(g)$ is still infinite.

The interested reader will be delighted to check that the "Extendable" and "Closure" hypotheses still hold at level $\alpha + 1$.

Now we turn to the limit case, which is the best part of the proof. Suppose that β is a limit ordinal, and that we've defined $\mathrm{Lev}_\alpha(T)$ for each $\alpha < \beta$ (this proof covers the case that $\beta = \omega$ as well). In light of the "Extendable" condition, we need to ensure that every f that we've constructed so far extends to a function defined on β. Thus let $f \in \bigcup_{\alpha<\beta} \mathrm{Lev}_\alpha(T)$ be given. Let α_0 so that $f \in \mathrm{Lev}_{\alpha_0}(T)$, and let $\alpha_0 < \alpha_1 < \alpha_2 < \dots$ be an increasing sequence of ordinals that is cofinal in β (meaning that every ξ below β is below one of the α_ns; the sequence goes "all the way up" to β). We will extend f to level α_1, then level α_2, and so on.

> Let's first try an instructive false start. By applying the "Extendable" hypothesis infinitely-many times, we can build f_1 on level α_1 extending f, then f_2 on level α_2 extending f_1, and so on, until we have a sequence $f \sqsubset f_1 \sqsubset f_2, \dots$. We can stitch them together to get an injective function g defined on β. The problem is that there's no reason to believe that g satisfies the "Options" hypothesis at β. I.e., there's no reason to believe that $\mathrm{ran}(g)$ isn't just all of $\mathbb{N}$. This would ruin the construction of the tree, since the node g would then be a leaf: there's nowhere to go with g! So we need to be more careful.

What's missing from the instructive false start is a mechanism that ensures that when we build the extensions $f_1, f_2, \dots$ and take the limit g, we still have that $\mathbb{N} \setminus \mathrm{ran}(g)$ is infinite. We will do this by also constructing an infinite set Z of natural numbers which we commit to keeping out of the ranges of the f_i. *But* since we can't commit to a given set Z of excluded values at the *start* of the construction, we have to also build this set Z by recursion *at the same time as we build the* f_i. We will, in fact, add one new value to the set Z each time we build a new f_i.

Thus to extend f to a function g defined on β, we will recursively construct functions f_i on level α_i of T, as well as **distinct** natural numbers $\mathrm{ctt}(1), \mathrm{ctt}(2), \mathrm{ctt}(3), \dots$ (here "ctt" stands for "can't touch this"[5]) so that for each $k \in \omega$, f_k is on level α_k of T, and

$$\{\mathrm{ctt}(1), \dots, \mathrm{ctt}(k)\} \cap \mathrm{ran}(f_k) = \emptyset.$$

Start by defining $f_0 = f$. Suppose that we've built f_k and defined the previous ctt values. Let g_{k+1} be an extension of f_k to level α_{k+1}. If $\mathrm{ran}(g_{k+1}) \cap \{\mathrm{ctt}(1), \dots, \mathrm{ctt}(k)\} = \emptyset$, then we can set $f_{k+1} = g_{k+1}$ and pick $\mathrm{ctt}(k+1)$ from $\mathbb{N} \setminus (\mathrm{ran}(g_{k+1} \cup \{\mathrm{ctt}(1), \dots, \mathrm{ctt}(k)\})$. Suppose, however, that $\mathrm{ran}(g_{k+1})$ does include some of the previous ctt values, say

$$b = \mathrm{ran}(g_{k+1}) \cap \{\mathrm{ctt}(1), \dots, \mathrm{ctt}(k)\} \neq \emptyset.$$

[5]Hot take: math papers should include more references to M.C. Hammer.

We perform some surgery on g_{k+1} to construct a new function f_{k+1} whose range has none of the previous ctt values.

By assumption, g_{k+1} is injective. Thus for each $l \in b$, there is a unique γ_l so that $g_{k+1}(\gamma_l) = l$. Since b is finite, the set of the γ_l is also finite. We will define a new function f_{k+1} which modifies g_{k+1} on precisely these γ_l values. Namely, set $f_{k+1}(\delta) = g_{k+1}(\delta)$ for each $\delta \notin \{\gamma_l : l \in b\}$. For the values on the γ_l, note that

$$\mathbb{N} \setminus (\mathrm{ran}(g_{k+1}) \cup \{\mathrm{ctt}(1), \dots, \mathrm{ctt}(k)\})$$

is infinite. Thus we can let $f_{k+1}(\gamma_l)$ select values from this set, while maintaining injectivity. Crucially, since f_{k+1} is a finite modification of g, the "Closure" hypothesis implies that f_{k+1} is on level α_{k+1} of the tree. This completes the definition of f_{k+1}. Finally, define $\mathrm{ctt}(k+1)$ to be any value in $\mathbb{N}$ that is not in $\mathrm{ran}(f_{k+1})$ and not equal to any previous ctt value.

This completes the construction of the sequences $f = f_0 \sqsubset f_1 \sqsubset f_2, \dots$ and $\mathrm{ctt}(1), \mathrm{ctt}(2), \dots$. By construction, the union of the f_k is a function f^* with domain β. f^* is injective, as it is the union of an increasing sequence of injective functions. Finally, f^* satisfies the "Options" hypothesis, since $\mathrm{ran}(f^*) \cap \{\mathrm{ctt}(k) : k \in \mathbb{N}\} = \emptyset$ and since $\{\mathrm{ctt}(k) : k \in \mathbb{N}\}$ is infinite.

We're almost done! Since f was arbitrary, we've now shown how to extend every function on a level below β to a function on β. Since there are only countably-many functions in T on levels below β, we've only added countably-many functions to β. To secure the "closure" condition, we also add every injective $h : \beta \to \mathbb{N}$ which is a finite modification of a function that we've already constructed. This also adds at most a countable infinity of values to level β. This completes the limit case of the construction of T, and with that, we've constructed T! □

It turns out that with a little more work, one can add a further interesting structural feature to the above construction. Namely, you can modify the above construction so that the resulting tree is a countable union of very simple pieces called **antichains**. The idea of decomposing a complicated object into simple parts is exactly what Ramsey theory is about! So we have a nice tie-in to the previous paper ([Gil25]) in this series.

So first, we will give the definition of an antichain. Then we will look at some examples from finite trees to get a feeling for what we're talking about. After that, we'll explore the additional structural feature that you can add to the construction in Theorem 5.5.

Definition 5.6. Let $(P, \leq)$ be a partially ordered set. An *antichain* is a subset $A \subset P$ so that for any two $s \neq t$ from A, neither $s \leq t$ nor $t \leq s$.

Thus in an antichain, no two elements can be compared: pick any two distinct elements from the antichain, and neither one is above the other.

Let's look at an example from the following finite tree.

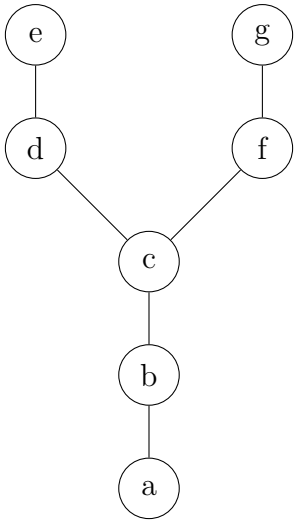

Recalling how we define the ordering in a finite tree, we see that $a < b < c$, that $c < d < e$, and that $c < f < g$. Note that by transitivity, we also have, for example, that $a < e$. But notice that $\{d, f\}$ is an antichain: neither d nor f are above each other. $\{d, g\}$ is another antichain, and so are $\{e, f\}$ and $\{e, g\}$. Notice that antichains are not the same as independent sets: the set $\{a, c, e, g\}$ is an independent set, since there are no edges between any two nodes from that set. But $\{a, c, e, g\}$ is most certainly not an antichain, since, for example, a and e are in the set, but $a < e$.

Thus we see that antichains in the context of the ordering on a tree are examples of independent sets in the context of graph theory, but the other direction definitely doesn't hold.

Let's now tie this in to the improved version of the construction in Theorem 5.5. We'll state the theorem, but skip the proof, for the sake of space:

Theorem 5.7. *There exist an ω_1-Aronszajn tree $(T, <)$ and a function $f : T \to \mathbb{Q}$ with the property that if $s < t$ are nodes in T, then $f(s) < f(t)$.*

In other words, you can also construct an order-preserving map from the tree into the rationals. Here's the connection to antichains and the reason why this is interesting: for each $q \in \mathbb{Q}$, consider the fiber $f^{-1}(\{q\}) = \{s \in T : f(s) = q\}$. This set is an antichain: if you take $s \neq r$ from T so that $f(s) = q = f(r)$, then neither $s < r$ nor $r < s$. The two elements, in other words, cannot be compared in the tree

ordering. Thus the set $f^{-1}(\{q\})$ is a very simple subset of T, namely, an antichain! And as we do this pullback for each $q \in \mathbb{Q}$, we cover all of the nodes in T, i.e.,

$$T = \bigcup_{q \in \mathbb{Q}} f^{-1}(\{q\}).$$

Thus T *is a union of countably-many antichains!* But since T is uncountable, and since a countable union of countable sets is countable, at least one of these sets in the union must be uncountable. Thus T **must have an uncountable antichain** (this will be important in the next section). This is therefore a kind of Ramsey-theoretic principle: we took a complicated uncountable object, $(T, <)$, and we decomposed it into countably-many simple parts, and by the pigeonhole principle, at least one of those simple parts must be uncountable.

Such trees have a name:

Definition 5.8. An ω_1-tree $(T, <)$ is called a *special ω_1-Aronszajn tree* if there exists a countable set Z and a function $f : T \to Z$ so that if $s < t$ are nodes in T, then $f(s) \neq f(t)$. Such an f is called a *specializing function*.

The cool kids say that such an f is "injective on chains."

Remark 5.9. *Note that if $(T, <)$ is an ω_1-tree and $f : T \to Z$ is a specializing function as in Definition 5.8, then $(T, <)$ is an Aronszajn tree: if $b \subset T$ is a chain, then as f is injective on chains, b must be countably infinite (any set which injects into a countable set is itself countable). Thus there are no cofinal branches through T, since any cofinal branch is an* uncountable *chain.*

We have one last result in this section before moving on. This concerns whether Aronszajn trees exist at the next cardinal above ω_1, namely ω_2. This question turns out to have a complicated but fascinating answer. For now, we'll sketch the proof that if the Continuum Hypothesis (CH) is true, then there is an ω_2-Aronszajn tree (namely, a tree of height ω_2, all of whose levels have size $\leq \omega_1$, but with no cofinal branch).

Theorem 5.10. *The CH implies that there is an ω_2-Aronszajn tree.*

Proof. We only sketch a few of the details, because it is almost identical to the proof that there is an ω_1-Aronszajn tree!

We construct T recursively. For an ordinal $\alpha < \omega_2$, $\mathrm{Lev}_\alpha(T)$ will consist of certain injective functions $f : \alpha \to \omega_1$ so that $\omega_1 \setminus \mathrm{ran}(f)$ is uncountable; this modifies the "Options" recursion hypothesis from the proof of Theorem 5.5. Here's how we modify the "Closure" recursion hypothesis from Theorem 5.5: suppose that $f \in \mathrm{Lev}_\alpha(T)$,

that $g : \alpha \to \omega_1$ is injective, and that $\{\xi < \alpha : f(\xi) \neq g(\xi)\}$ is at most countable. Then $g \in \mathrm{Lev}_\alpha(T)$. You should be worried that this makes $\mathrm{Lev}_\alpha(T)$ too big: if $f : \alpha \to \omega_1$, where $\alpha < \omega_2$, then there are $\omega_1^\omega = 2^\omega$-many functions $g : \alpha \to \omega_1$ so that $\{\xi < \alpha : f(\xi) \neq g(\xi)\}$ is at most countable. But under the CH, 2^ω is equal to ω_1! Thus there are only ω_1-many such g, which in turn implies that the levels of T will not have cardinality exceeding ω_1 (remember that we have to make sure that T isn't too wide).

Other than this observation, the proof is nearly identical to the proof of Theorem 5.5. Give it a try! □

6. Suslin's problem and the uniqueness of $\mathbb{R}$

We now turn to a connection between ω_1-trees and properties of the real numbers. A crucial fact about $\mathbb{R}$ is that $\mathbb{R}$ is the unique complete, ordered field. This means that if $\mathbb{F}$ is any complete, ordered field then there is an isomorphism between $\mathbb{F}$ and $\mathbb{R}$.[6] We call this a *characterization* of $\mathbb{R}$.

6.1. Suslin trees and a question about $\mathbb{R}$

Let's now start playing around with this characterization. We begin with a standard fact about $\mathbb{R}$, namely that the set of rational numbers, $\mathbb{Q}$, is a dense subset of $\mathbb{R}$. This means that if $a < b$ are real numbers, then there exists a rational $q \in \mathbb{Q}$ so that $a < q < b$. Having a countable, dense subset is sufficiently important that it gets its own definition:

Definition 6.1. A linear order $(L, <)$ is *separable* if it has a countable, dense subset.

Note that we can also rephrase the uniqueness of $\mathbb{R}$ as follows: $\mathbb{R}$ is the unique complete, dense, unbounded, linear order which is also separable.

Remark 6.2. *Make sure you get the main idea: we have a characterization of* $\mathbb{R}$*, and it has something to do (among other things) with the fact that the* **countable** *set* $\mathbb{Q}$ *is dense.*

A corollary of the density and countability of $\mathbb{Q}$ (i.e., a corollary of the fact that $\mathbb{Q}$ witnesses the separability of $\mathbb{R}$) is this: if $(I_\lambda)_{\lambda \in \Lambda}$ is a sequence of non-empty, pairwise disjoint, open intervals in $\mathbb{R}$, then the set $\{I_\lambda : \lambda \in \Lambda\}$ is countable. Here's the proof: for each $\lambda \in \Lambda$, use the density of $\mathbb{Q}$ to find a rational $q_\lambda \in I_\lambda$. Now observe that

[6] Dedikind's work in [Ded72] laid the groundwork for this characterization of $\mathbb{R}$ as the unique complete, ordered field.

the function $(\lambda \in \Lambda) \mapsto q_\lambda$ is an injection: if $\lambda \neq \lambda'$ are in Λ, then I_λ and $I_{\lambda'}$ are disjoint, by assumption. Thus $q_\lambda \neq q_{\lambda'}$. Consequently, Λ is a countable set, seeing that it injects into $\mathbb{Q}$ (any set which injects into a countable set is itself countable).

To summarize, $\mathbb{R}$ satisfies that any set of pairwise disjoint, non-empty open intervals is countable, and we proved this using the density of the countable set $\mathbb{Q}$. This property is also important enough to have its own definition:

Definition 6.3. Let $(L, <)$ be a dense, linear order. $(L, <)$ satisfies the *countable chain condition (c.c.c.)* if any set of pairwise disjoint, non-empty open intervals in $(L, <)$ is countable.

In general, it's true that **every separable linear order satisfies the c.c.c.** Suslin essentially asked whether we can replace "separable" with "c.c.c." in the characterization of $\mathbb{R}$. More precisely, consider the following question:

Question 6.4. Suppose that $(L, <)$ is a complete, dense, unbounded, linear order that satisfies the **countable chain condition**. Is $(L, <)$ isomorphic to the reals, i.e., is $(L, <)$ separable?

This question is part and parcel of what we do in math: after proving a theorem, we want to see if we can weaken the assumptions and still get as strong of a result.

Remark 6.5. *Make sure you still have the narrative: we've characterized* $\mathbb{R}$ *using separability (among other things), and we'd like to know if a certain consequence of separability, namely the c.c.c., is strong enough to uniquely characterize* $\mathbb{R}$*.*

Here it is helpful to have some more terminology:

Definition 6.6. A *Suslin line* is a complete, dense, unbounded, linearly ordered set that has the countable chain condition but which is **not** separable.

Thus if a Suslin line exists, then the answer to Question 6.4 is negative, and if there are no Suslin lines, then Question 6.4 has a positive answer. But it turns out that this question of whether or not a Suslin line exists cannot be decided on the basis of the usual axioms of ZFC set theory![7] This means that if ZFC is consistent, then so are the following two theories:

1. ZFC+ "there is a Suslin line" and
2. ZFC+ "no Suslin line exists".

[7]It's as if this question looks to the axioms of ZFC and says, "can't touch this."

Given the constraints of space (and in order to show mercy on the reader) we're only going to focus on item (1) above. Item (2) would take us into very different territory.[8]

One of the first things to say about (1) is that it turns out to be a little easier to do the combinatorics with trees, rather than with linear orders. And it also turns out that having a Suslin line is equivalent to having a certain type of ω_1-tree. Here's the definition of the latter tree (which we'll connect to a Suslin line in a bit):

Definition 6.7. An *ω_1-Suslin tree* is an ω_1-tree $(T, <)$ so that every antichain in $(T, <)$ is countable.

Remark 6.8. *Keep track of the narrative! We're connecting the question of whether a certain list of properties (among which is the c.c.c.) uniquely characterizes $\mathbb{R}$ to the question of whether a certain kind of ω_1-tree exists. However, we can't decide which way the question goes just on the basis of the* ZFC *axioms for set theory.*

Before stating the equivalence, let's connect these ideas to the ones in the previous section by proving that every ω_1-Suslin tree is ω_1-Aronszajn.

Lemma 6.9. *Every ω_1-Suslin tree is an ω_1-Aronszajn tree, but is not a special ω_1-Aronszajn tree.*

Proof. Let T be an ω_1-tree which is not Aronszajn, and we will show that it is not Suslin. By definition, we can choose a cofinal branch b through T, and since T has height ω_1, we can let b_α denote the node of b on level α, for each $\alpha < \omega_1$. We need to use this branch to create an uncountable antichain. We will do so by splitting off each node b_α: we have infinitely many nodes of T above b_α on level $\alpha + 1$, and we'll pick one that is different from $b_{\alpha+1}$. These "split off" nodes will then create an antichain.

In more detail, for each α, let $a_{\alpha+1}$ be a node on level $\alpha + 1$ of T which is above b_α and which is distinct from $b_{\alpha+1}$. Here's a picture to illustrate the idea:

[8]Namely, into iterated forcing. The proof that (2) is consistent (assuming that ZFC is), is due to Solovay and Tennenbaum in their watershed paper [ST71].

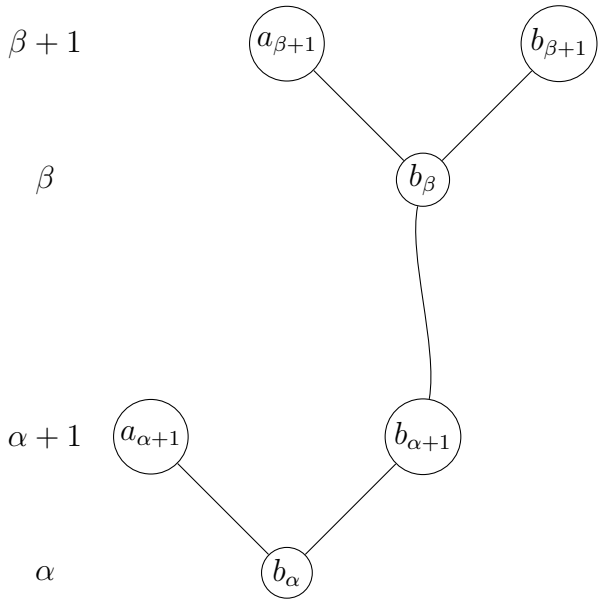

The ordinals in the left column $(\alpha, \alpha+1, \beta, \beta+1)$ represent levels in the tree, where $\alpha < \beta$. The curvy line from $b_{\alpha+1}$ to b_β represents the fact that $b_{\alpha+1}$ is below b_β, since they both come from the branch b.

We claim that $A := \{a_{\alpha+1} : \alpha < \omega_1\}$ is an uncountable antichain. It's uncountable since $\alpha \mapsto a_{\alpha+1}$ is an injection. Suppose for a contradiction that A is not an antichain. Pick $\alpha < \beta$ so that $a_{\alpha+1} < a_{\beta+1}$ in the tree ordering. By definition, $a_{\beta+1}$ is above b_β, which is in turn above $b_{\alpha+1}$. Moreover, by assumption, $a_{\alpha+1}$ is below $a_{\beta+1}$. Thus both $b_{\alpha+1}$ and $a_{\alpha+1}$ are below $a_{\beta+1}$. By definition of a tree, the set of predecessors of any node, and hence of $a_{\beta+1}$, is a well-order. Thus either $a_{\alpha+1} < b_{\alpha+1}$ or vice-versa. But this is impossible, as $a_{\alpha+1}$ and $b_{\alpha+1}$ are distinct nodes on the same level, and hence neither is above the other.[9]

For the second part of the lemma, recall from earlier that if T is special, then T is the union of countably-many antichains. Since T is uncountable, at least one of those antichains must be uncountable, and this contradicts the definition of Suslin. □

And now for the connection between Suslin lines and Suslin trees:

Theorem 6.10. *A Suslin line exists iff an ω_1-Suslin tree exists.*

Proof. Here we only give a very rough sketch of the proof.

[9]Another, somewhat loose, way to look at it is that if $a_{\alpha+1}$ is below $a_{\beta+1}$, then you get a loop-like thing in the above picture, which is impossible in a tree. The issue with this is that things like "loop" don't make as much sense for trees in the context of higher infinities. Indeed, it's the goal of this proof to make this loose idea precise!

Suppose that you're given a Suslin line $(L, <)$. Then you can construct a tree where the nodes are closed intervals $[a, b]$ (with $a < b$) from $(L, <)$, and where, given two such intervals I, J, we say that $I \leq J$ iff $J \subseteq I$ (so as you go higher up on the tree, the intervals shrink). The fact that $(L, <)$ is not separable ensures that the tree has height ω_1, and the fact that $(L, <)$ has the c.c.c. ensures that the tree has neither uncountable branches nor uncountable antichains.

Given a Suslin tree $(T, <)$, you can define a linear order on the branches of T, roughly by ordering them lexicographically. Using the assumed properties of the tree T, one can deduce that $(L, <)$ is a Suslin line.[10] □

Thus, Question 6.4 has a negative answer iff there is an ω_1-Suslin tree! This is satisfying to set theorists, since trees are a little more amenable to set-theoretic techniques.

6.2. **A sufficient condition for constructing a Suslin tree**

We will now describe how one proves that it's consistent with the ZFC axioms of set theory that there is a Suslin line (equivalently, tree).

Constructing a Suslin tree gives us an excuse to talk about another lovely bit of combinatorics, namely, diamond principles.[11] Roughly speaking, a diamond principle is a kind of "guessing" principle. On ω_1, the most well-known diamond principle says that there is a sequence of "guesses" $\langle A_\alpha : \alpha < \omega_1 \rangle$ with $A_\alpha \subseteq \alpha$, so that for *every* $A \subseteq \omega_1$, there are plenty[12] of α so that $A \cap \alpha = A_\alpha$. It's a deja vu thing: you see lots of the initial segments (approximations) of A living on the sequence.

Typically, in the context of a diamond principle, the set A represents some "obstacle" that must be dealt with at the *end* of a given combinatorial construction. And you can argue that you deal with the full, global obstacle A by ensuring that you deal with enough of the approximations A_α throughout the course of the construction. In other words, a $\diamondsuit$ principle often ensures that you can solve a global problem at the end of a construction by solving enough local, similar problems throughout the course of the construction.

Now here's the official definition:

Definition 6.11. A *Diamond Sequence*, or $\diamondsuit$, is a sequence $\langle A_\alpha : \alpha < \omega_1 \rangle$ where

[10] There are a few more things that you need to assume about the tree, but you pretty much get these for free. See Chapter 9 of Jech's textbook [Jec03].

[11] This is only the tip of the iceberg, using what are now pretty old techniques. For some of the cutting-edge research, see the excellent papers by Brodsky and Rinot: [BR17], [BR21].

[12] The precise notion is "stationary", which is a version of positive measure with respect to the so-called club filter on ω_1.

1. for each $\alpha < \omega_1$, $A_\alpha \subseteq \alpha$, and
2. for each $A \subseteq \omega_1$, there are "plenty" of α so that $A \cap \alpha = A_\alpha$.

Now here are two crucial facts about $\diamondsuit$:

Proposition 6.12.

1. *If there is a $\diamondsuit$ sequence, then the* CH *holds (i.e., there are no infinite sizes that sit between countable infinity and the size of the continuum).*
2. *In Gödel's constructible universe, denoted L, there is a $\diamondsuit$ sequence.*[13] *(This result is due to Jensen in his landmark paper [Jen72].)*

Recall that if ZFC is consistent, so is ZFC and the statement that every set is in Gödel's L (see Gödel's monograph [Göd40]). And since L contains $\diamondsuit$ sequences, it is consistent that there are $\diamondsuit$ sequences.

Remark 6.13. *To summarize: we're investigating a certain kind of combinatorial principle, $\diamondsuit$, because we know that it's consistent that $\diamondsuit$ sequences exist, and we know that if we have a $\diamondsuit$ sequence, we can construct a Suslin tree. And from a Suslin tree, we can construct a Suslin line!*

Now we sketch the construction of a Suslin tree from a $\diamondsuit$ sequence; our goal is to convey the main ideas but without going through all of the (lovely) details.

Theorem 6.14. *If there is a $\diamondsuit$ sequence, then there is an ω_1-Suslin tree.*

Proof. We will only sketch the idea of the proof here, highlighting the role that $\diamondsuit$ plays.

Like most constructions of set-theoretic trees, we construct the ω_1-Suslin tree $(T, <)$ by recursion, going level-by-level. Make sure to keep the goal in mind: we have to build T in such a way that, at the end of the construction, there are no uncountable antichains. This means that, given any uncountable set $A \subseteq T$ of nodes, we need to be able to find two nodes $x, y \in A$ with either $x < y$ or $y < x$.

What's difficult about this is that it's only *after* constructing $(T, <)$ that "the enemy" gives us their favorite $A \subseteq T$.

The way around this: while we can't "see" an uncountable $A \subseteq T$ when we're at a local level in the construction (because we don't know the full tree yet!), we *can* see the "guesses" from the $\diamondsuit$ sequence.

[13] I once had a buddy who gave a talk on this with the title "L in the sky with diamonds." It's a joke about the Beatles.

And here's why this is good news: at limit stages $\alpha < \omega_1$, we will consider the set A_α (which could potentially equal $A \cap \alpha$ for one of the A's that the enemy has in their back pocket). We will build the next level of the tree in such a way that A_α, if it is a maximal antichain in the tree up to level α (and since $A_\alpha \subseteq \alpha$ and α is countable, A_α must be a countable antichain), does not grow into an uncountable antichain. We will accomplish this by ensuring that A_α stays maximal as we go about the rest of the tree construction. We call this ***sealing off*** the antichain.

In a bit more detail: suppose that $\alpha < \omega_1$ is a limit ordinal and that we've built $\mathrm{Lev}_\xi(T)$ for every $\xi < \alpha$. Consider A_α as a subset of $\bigcup_{\xi<\alpha} \mathrm{Lev}_\xi(T)$ (i.e., as a subset of what we've built so far). If A_α is not an antichain in this tree, then we don't need to do anything special at this stage, and we can continue as in the proof of Theorem 5.5. On the other hand, suppose that A_α is an antichain in the portion of the tree that we've constructed so far. We may assume, without loss of generality, that A_α is a maximal antichain. Using this maximality, we can make sure that when we build level α (i.e., the next level we have to build) every node x on level α extends a node $z \in A_\alpha$. *This in turn ensures that every node on every level $\beta \geq \alpha$ will extend a node of A_α, and from this we can prove that A_α is still maximal in the whole tree!* Here's a picture:

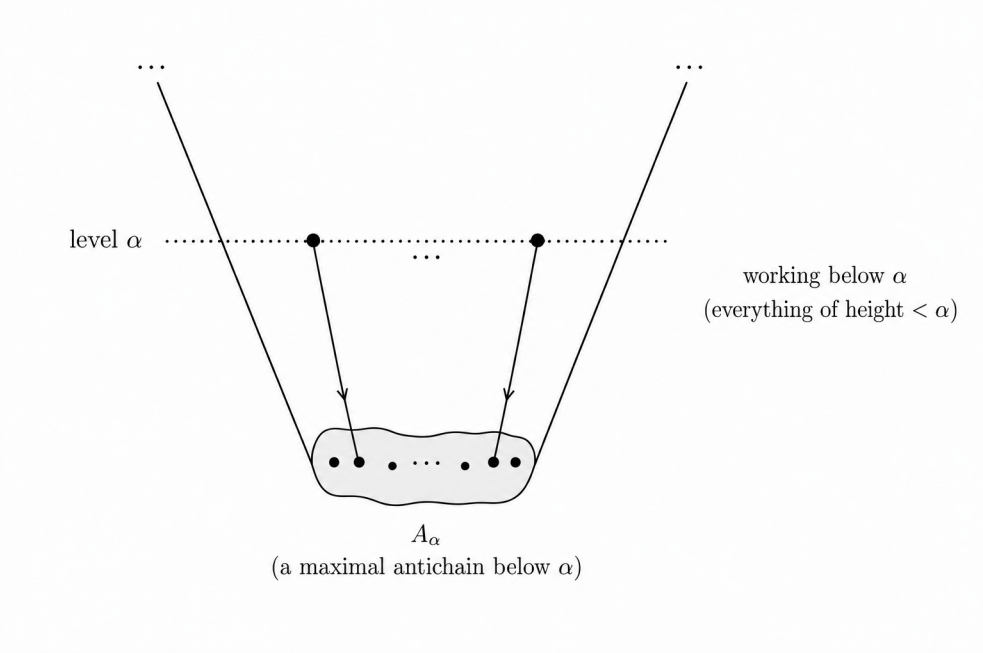

Let's emphasize the order of the quantifiers: for each node x on level α, we find a node of A_α to put below it (this order of quantifiers is illustrated in the picture by the arrows going *from* level α nodes *to* nodes in A_α). Thus no matter what we do after this stage in the construction, every node on level α or higher is above a node in

A_α, and this is enough to ensure that A_α remains a maximal antichain *in the whole tree*. A_α is therefore "sealed off."

And here's how this all ties together: suppose we've constructed $(T, <)$ according to this recipe. We now argue that every antichain in T is countable. Thus fix $A \subseteq T$, an antichain. Extend A if necessary to make it a maximal antichain. Using the definition of our $\diamondsuit$ sequence, we may find a limit ordinal α so that $A \cap \alpha = A_\alpha$ and so that A_α is maximal in $\bigcup_{\xi<\alpha} \mathrm{Lev}_\xi(T)$. By construction, we know that A_α (which, again, equals $A \cap \alpha$) remains a maximal antichain in the whole tree. But $A \cap \alpha = A_\alpha$ is maximal in the whole tree, and $A_\alpha = A \cap \alpha \subseteq A$, while A is also maximal in the whole tree. Since a maximal antichain can't get any bigger (while remaining an antichain), we must have that $A_\alpha = A$, or in other words, that $A = A \cap \alpha$! Consequently, A is countable. □

Thus we see that it is consistent that an ω_1-Suslin tree exists (since we can work in Gödel's L), and hence that a Suslin line exists.

7. **Summary and acknowledgments**

So what have we seen in this paper? We started by discussing graph-theoretic trees and then, using the theory of well-orders, we generalized this to much "higher" trees. We saw that König's result (Theorem 2.3) does *not* generalize to the uncountable (in the most straightforward way) by constructing ω_1-Aronszajn trees. Then we discussed ω_1-Suslin trees and their connection to properties of the real numbers. I hope that this has whet your appetite for infinitary combinatorics!

I'd like to thank Leonardo Finzi for the invitation to submit a part 2 in the series.

I used AI tools to generate the images and figures in the paper.

References

[BR17] Ari Meir Brodsky and Assaf Rinot. A microscopic approach to souslin-tree constructions, part i. *Annals of Pure and Applied Logic*, 168(11):1949–2007, 2017.

[BR21] Ari Meir Brodsky and Assaf Rinot. A microscopic approach to souslin-tree construction, part ii. *Annals of Pure and Applied Logic*, 172(5):102904, 2021.

[Ded72] Richard Dedekind. *Stetigkeit und irrationale Zahlen*. Friedrich Vieweg und Sohn, Braunschweig, 1872.

[Gil25] Thomas Gilton. When a set theorist hears "combinatorics": Infinite Ramsey theory. *Pittsburgh Interdisciplinary Mathematics Review*, 3:28–52, 2025.

[Göd40] Kurt Gödel. *The Consistency of the Continuum Hypothesis*, volume 3 of *Annals of Mathematics Studies*. Princeton University Press, Princeton, NJ, 1940.

[Jec03] Thomas Jech. *Set Theory: The Third Millennium Edition, Revised and Expanded*. Springer Monographs in Mathematics. Springer, Berlin, Heidelberg, 2003.

[Jen72] Ronald Björn Jensen. The fine structure of the constructible hierarchy. *Annals of Mathematical Logic*, 4(3):229–308, 1972.

[Kön27] Dénes König. Über eine schlussweise aus dem endlichen ins unendliche. *Acta Litterarum Academiae Scientiarum Hungaricae (Szeged)*, 3:121–130, 1927.

[Kur35] Dimitrije Kurepa. Ensembles ordonnés et ramifiés. *Publications Mathématiques de l'Université de Belgrade*, 4:1–138, 1935.

[ST71] Robert M. Solovay and Stanley Tennenbaum. Iterated cohen extensions and souslin's problems. *Annals of Mathematics*, 94(2):201–245, 1971.

Thomas Gilton
Department of Mathematics, University of Pittsburgh, Pittsburgh, PA
Email: tdg25@pitt.edu